A Survey of AMR Applications

Shira Wein
Amherst College
swein@amherst.edu

Juri Opitz
University of Zurich
opitz.sci@gmail.com

Abstract

In the ten years since the development of the Abstract Meaning Representation (AMR) formalism, substantial progress has been made on AMR-related tasks such as parsing and alignment. Still, the engineering applications of AMR are not fully understood. In this survey, we categorize and characterize more than 100 papers which use AMR for downstream tasks—the first survey of this kind for AMR. Specifically, we highlight (1) the range of applications for which AMR has been harnessed, and (2) the techniques for incorporating AMR into those applications. We also detect broader AMR engineering patterns and outline areas of future work that seem ripe for AMR incorporation. We hope that this survey will be useful to those interested in using AMR and that it sparks discussion on the role of symbolic representations in the age of neural-focused NLP research.

1 Introduction

Abstract Meaning Representation (AMR; Banarescu et al., 2013) is a semantic representation that takes the form of a rooted, directed graph. Since the release of AMR in 2013, a full AMR-ecosystem has emerged, with substantial research activity on AMR annotation, text-to-AMR parsing, AMR-to-text generation, and domain- and language-based extensions of AMR.[1] In particular, the progress on text-to-AMR parsing and AMR-to-text generation has propelled work using AMR for various NLP applications. To date, downstream applications of AMR have been spread across numerous tasks and have found varying degrees of success.

Thus, given the recent advancements for and with AMR, this survey addresses the pressing question: how can AMR be used for engineering purposes and downstream applications? Our main goals of this investigation include (1) providing an overview of the many *application areas and tasks* where AMR has been applied, (2) examining *what techniques* have been used to leverage AMR for NLP systems, and (3) detecting *new avenues* for future applications of AMR in NLP research.

Our investigation is also motivated by the prevalence of large language models (LLMs) that seem to be able to generalize across a large suite of NLP tasks, prompting consideration of how semantic representations can remain useful. We hope that

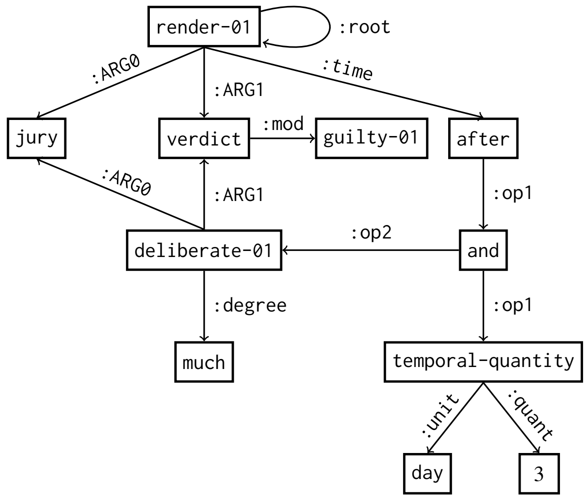


```
(r / render-01
    :ARG0 (j / jury)
    :time (a / after
        :op1 (a2 / and
            :op1 (t / temporal-quantity
                :unit (d / day)
                :quant 3)
            :op2 (d2 / deliberate-01
                :degree (m / much)
                :ARG0 j
                :ARG1 v)))
    :ARG1 (v / verdict
        :mod (g / guilty-01)))
```

Figure 1: The AMR for the sentence "After 3 days and much deliberation, the jury rendered a guilty verdict," as a graph (top) and as a string in PENMAN notation (bottom).

[1]Currently, the AMR Bibliography contains more than 450 papers: https://nert-nlp.github.io/AMR-Bibliography/.

Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, pages 6856–6875
November 12-16, 2024 ©2024 Association for Computational Linguistics

our survey will serve both as a useful starting point and a source of inspiration to those interested in working with AMR.

2 Abstract Meaning Representation

2.1 AMR Formalism

Semantic representations such as AMR aim to convey the meaning of a text and can be designed to focus on specific aspects of meaning. AMR is specifically designed to reflect "who does what to whom" as the schema centers on predicate-argument relations. By abstracting away from the surface form, two sentences with equivalent meaning and content words should be represented by the same AMR graph. Among semantic representations, AMR is particularly popular and well-resourced (Sadeddine et al., 2024).

AMRs are rooted and directed, and can be represented in graph form or in the text-based PENMAN notation (Kasper, 1989) (an example AMR in both forms is shown in Figure 1); text-based AMRs are also called *linearized*, and often appear condensed onto one line for ease of neural encoding. Concepts correspond to nodes in the graph, and edges denote relationships between those concepts. These concepts can occupy core argument roles (i.e. `:argN`) or non-core roles (e.g. `:time` and `:domain`). AMR makes use of PropBank frame files (Palmer et al., 2005) to indicate the sense of each concept in the graph, as well as to specify the arguments associated with each concept. AMR annotation is unanchored, so individual tokens do not necessarily align with specific concepts in the graph. Coreferent concepts are reflected in AMR graphs as re-entrant graph nodes.

AMR also notably does not represent morphology or tense, meaning that annotation is fairly lightweight. Inter-annotator agreement is typically measured quantitatively using Smatch (Cai and Knight, 2013), which calculates graph overlap via hill climbing. Embedding-based metrics which measure AMR graph overlap include monolingual S2match (Opitz et al., 2020) and multilingual XS2match (Wein and Schneider, 2022).

AMR was originally designed for English and was not intended to serve as an interlingua (Banarescu et al., 2013), but the schema has since been considered for or adapted to numerous other languages: Czech (Urešová et al., 2014), Chinese (Xue et al., 2014; Li et al., 2016), Spanish (Migueles-Abraira et al., 2018; Wein et al., 2022), Vietnamese (Linh and Nguyen, 2019), Korean (Choe et al., 2020), Portuguese (Sobrevilla Cabezudo and Pardo, 2019; Anchiêta and Pardo, 2018; Inácio et al., 2022; Baptista et al., 2024), Turkish (Azin and Eryiğit, 2019; Oral et al., 2022), Persian (Takhshid et al., 2022), and German (Otto et al., 2024).

Multilingual adaptations of AMR which are not specific to one individual language include Uniform Meaning Representation (UMR; Van Gysel et al., 2021) and BabelNet (Martínez Lorenzo et al., 2022). Some extensions of the AMR schema incorporate tense and aspect (Donatelli et al., 2018; Bakal, 2021), while others move beyond the sentence level (O'Gorman et al., 2018; Moreda et al., 2018; Naseem et al., 2022). Many engineering applications of AMR have focused on English, likely due to English AMR tools currently being the most widely available and accurate.

2.2 AMR Parsing and Text Generation

Two crucial AMR-intrinsic tasks are text-to-AMR **parsing** and AMR-to-text **generation**. Both tasks are actively researched, monolingually and multilingually. Substantial efforts towards highly accurate parsing and generation contribute further to the interest in using AMR for downstream applications.

Parsing and generation models now tend to leverage pre-trained Transformers that are fine-tuned on linearized AMR graphs (Bevilacqua et al., 2021). Current models can parse and generate quite accurately, reporting Smatch scores upwards of 86% for parsing (Lee et al., 2022b; Vasylenko et al., 2023) and more than 50 BLEU (Papineni et al., 2002) points for generation (Cheng et al., 2022).

Thus, while AMR parsing and generation are not yet solved (Opitz and Frank, 2022a; Groschwitz et al., 2023), model performance is quite high, and success towards semantically consistent parsing and generation (Kachwala et al., 2024) has led to a spike in the downstream utility of AMR.

3 Applications of AMR

In this section, we will discuss the broad categories of downstream uses of AMR in natural language processing. We present more than 20 individual areas of application, and discuss how each work incorporated AMR.

3.1 AMR for Meaning-Focused Tasks

Intuitively, the tasks which have most often seen AMR incorporated are *tasks which focus on elu-*

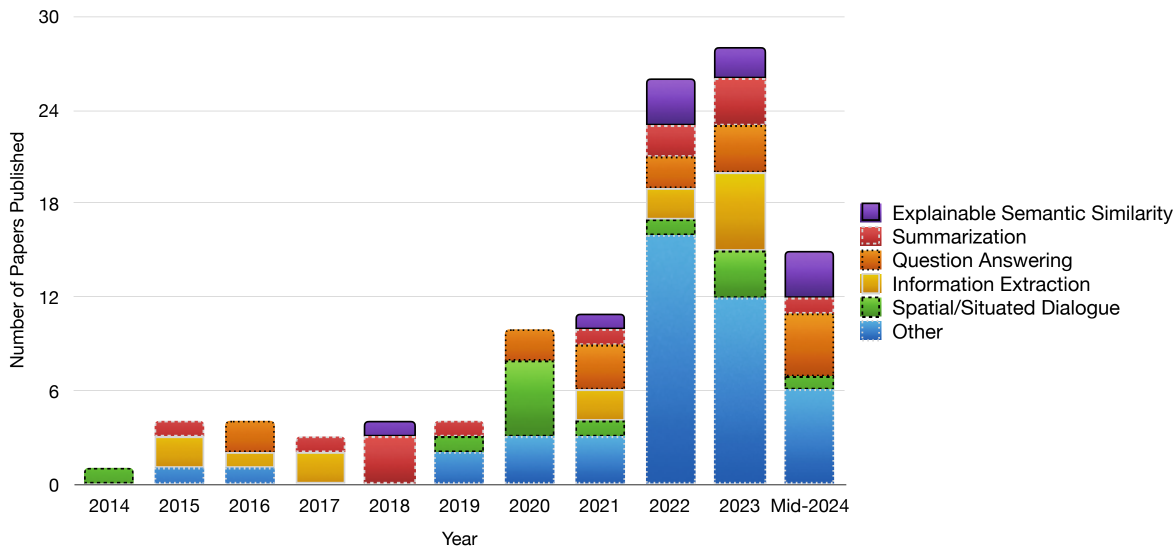

Figure 2: Bar chart of the number of papers using AMR in downstream applications per year, from 2014 to 2024 (year to date). The 5 most common application areas are individually shown, with all other areas grouped into the "Other" category.

cidating the core elements of meaning. Broadly, the meaning-focused tasks which have seen AMR leveraged fall under information extraction, question answering, and summarization; these areas overlap, particularly information retrieval and question answering, as the former can be an important step for the latter.

Information Retrieval/Extraction. Early AMR investigations for information retrieval/extraction focused on the **biomedical domain**. Biomolecular interactions elicited via AMR have been used in classifiers to outperform state-of-the-art interaction models (Garg et al., 2016; Wang et al., 2017). Notably, Rao et al. (2017) showed that biomedical events are subgraphs of full AMR graphs and developed an LSTM model (Hochreiter and Schmidhuber, 1997) to identify those event subgraphs. Zhang et al. (2021) performed biomedical information extraction by creating an AMR graph enhanced with information from an external knowledge base. These enhanced AMR graphs were then encoded into a graph attention network, leading to an improvement over state-of-the-art methods.

Models for **event extraction** have also incorporated AMR, occasionally outperforming state-of-the-art models. First, Li et al. (2015) added AMR features in the form of node-relation unigrams and bigrams to an event detection model. More recently, Xu et al. (2023) created new event extraction labels by using an existing event extraction model and an AMR parser to compute a compatibility score between the event and an argument. Yang et al. (2023) performed event structure extraction by identifying whether there is an edge connecting the event and the argument, parsing an AMR and then using a Graph Neural Network (GNN) to predict whether there is an edge. Again outperforming state-of-the-art event extraction models, Hsu et al. (2023) produced a linearized AMR, encoded it with a neural network, and prepended the encoding to the neural text embedding.

More specifically than general event extraction, Zhang and Ji (2021) performed **entity and relation extraction** by training an AMR encoder and using AMR parses in order to determine the order of decoded events. This work outperformed prior state-of-the-art models on information extraction by multiple F1 points. Gururaja et al. (2023) compared the utility of different sorts of linguistic graphs for Transformer-based models for relation extraction, finding that AMRs were most useful in few-shot settings. Pan et al. (2015); Steinmetz (2023) performed entity linking by mapping named entities onto concepts in AMR graphs.

At the **document level**, Xu et al. (2022) performed event extraction using text embeddings combined with document-level and sentence-level AMR graphs, outperforming prior state-of-the-art systems, and Zhao et al. (2023) encoded document-level AMR graphs in a GNN for relation extraction.

Finally, Müller and Kuwertz (2022) extracted relevant information from remote sensing database management systems, using AMR graph overlap metrics to measure semantic relevance.

Question Answering and Knowledge Graphs. When incorporating AMRs into question answering models, prior approaches have combined AMR graphs with a formal reasoning layer (Mitra and Baral, 2016) and sentence embeddings (Park et al., 2024). On the other hand, Bonial et al. (2020b) used AMR graphs directly, parsing medical questions (about COVID-19) into AMRs and comparing them against AMR graphs of possible answers. The answers were then ranked by similarity and the most similar response AMR was returned as the answer.

Regan et al. (2024) created multilingual AMR graphs of questions and developed a joint AMR-SPARQL parsing model for **hallucination detection** in knowledge base question answering (KBQA). For direct use in **KBQA**, AMR graphs have been converted into SPARQL queries (Bornea et al., 2021; Kapanipathi et al., 2021; Shivashankar et al., 2022). Similarly, AMRs have been used to produce Resource Description Framework (RDF) knowledge graphs (Burns et al., 2016; Meloni et al., 2017; Gangemi et al., 2023), and to semantic roles for a climate-focused knowledge base (Islam et al., 2022).

For **multi-hop question answering** (questions which require multiple steps to reach the answer), Xu et al. (2021) parsed AMR graphs of the hypothesis and the relevant facts and merged them, while Deng et al. (2022) segmented AMR parses of the question into subgraphs, and generated subquestions via AMR-to-text generation of the subgraphs. Similarly, for **open domain question answering** (ODQA), Wang et al. (2023) integrated AMR graphs of the relevant facts from a text by appending a single token embedding of each concept or relation in the AMR graph to the text embedding. Shi et al. (2024) performed ODQA via retrieval augmented generation (RAG), using an AMR-based algorithm to compress textual information into individual concepts. Pham et al. (2024) conditioned QA systems on AMR graphs, finding the approach works best with small models, which then outperformed very large LLMs such as ChatGPT.

The task of **machine comprehension**, which involves systems producing answers about a text, has also benefited from comparison between text and AMR graphs (Galitsky, 2020), with Sachan and Xing (2016) framing machine comprehension as a graph entailment problem. Towards question answer **dataset creation**, Rakshit and Flanigan (2021) parsed AMRs of sentences to generate question-answer pairs.

Summarization. AMR use in summarization has taken various approaches, often by parsing and joining AMR graphs of the sentences determined to be the most important. For instance, Dohare et al. (2017); Liao et al. (2018) picked the most important sentences from a text and created a single AMR graph from those sentences, then generated a short summary from key subgraphs.

Rather than parsing only the most important sentences, Liu et al. (2015) parsed individual AMR graphs of a text, combining them into one “summarization” graph by collapsing multiple concepts into single nodes with new concept labels, and then generating text from the summarized AMR. Variations of this approach have been proposed by Hardy and Vlachos (2018); Kouris et al. (2022).

Numerous approaches to **genre-specific summarization**, such as opinion summarization (Inácio and Pardo, 2021), TV transcript summarization (Hua et al., 2022), timeline generation (Mansouri et al., 2023), long dialogue summarization (Hua et al., 2023), and abstractive summarization of biomedical documents (Frisoni et al., 2023) have all seen the incorporation of AMR. Some of these works have leveraged AMRs by parsing AMRs of the text and then incorporating them into an LLM via an attention mechanism (Hua et al., 2022, 2023; Frisoni et al., 2023). In a non-English setting, Severina and Khodra (2019) used AMR for Indonesian multi-document summarization.

As a **post-hoc refinement step** for text summarization, Ming et al. (2018) used AMR and WordNet (Miller, 1995; Fellbaum, 1998) to filter out redundant information.

3.2 AMR to Abstract Away from the Surface

Following the logic that AMR “abstracts” away from the surface form, recent work has exploited AMR to produce new text with the same meaning as the original text. This work often uses AMR as an intermediate representation, i.e., parses an AMR from a text and then generates new text from the parsed AMR.

Style Transfer. Jangra et al. (2022) leveraged AMR as an intermediate representation to generate a paraphrase in a different style, using a fine-tuned

Transformer-based AMR parser (encoder) and multiple Transformer-based text generators (decoders) for various text styles. Shi et al. (2023) also used AMR as an intermediary, parsing an AMR graph from the text and performing concept-level style rewriting on the AMR graph (modifying the words to a different genre), achieving state-of-the-art results.

Paraphrase Generation. Prior work has utilized AMR for paraphrase generation by producing paraphrases directly from AMR graphs (Huang et al., 2023; Bao et al., 2023), occasionally altered with additional information (Lee et al., 2022a; Tu et al., 2024), or by injecting an embedding of an AMR into a Transformer model (Huang et al., 2022).

As AMR inherently preserves semantic similarity, Huang et al. (2023) used AMR directly as an intermediary to generate syntactically diverse paraphrase sets. They changed the root of the parsed AMR in order to be able to produce multiple AMRs,[2] and thus multiple sentences via AMR-to-text generation. Similarly, Shou et al. (2022); Ghosh et al. (2024) tackled data augmentation by parsing AMR graphs of the text, then editing them and generating new text.

Grammatical Error Correction. Cao and Zhao (2023) constructed denoised AMR graphs of sentences with grammatical errors and incorporated them into a sequence-to-sequence model as additional knowledge, achieving significantly higher precision and recall than the text-only baseline.

Machine Translation. Multiple approaches to neural machine translation have seen performance improvements when incorporating AMR as additional knowledge (Song et al., 2019; Nguyen et al., 2021). Li and Flanigan (2022) in particular observed performance gains when integrating AMR graphs into both the encoder and decoder of a Transformer model. Jin et al. (2024) on the other hand found that feeding an AMR in a zero-shot prompting setting with LLMs did not improve—or even hurt—performance.

AMR can also be used as an intermediary for a translation post-processing step in order to reduce the presence of translation artifacts (*translationese*, Wein and Schneider (2024b)).

[2]In AMR, the root indicates the linguistic *focus* of a sentence. Thus, changing the root of the AMR of the sentence "the cat drinks water" from *drink* to *water*, will yield a paraphrase such as "it is water that the cat drinks."

3.3 AMR for Domain-Specific Adaptations

Given that some syntactic content is not included in AMR graphs, the AMR schema has been adapted as necessary for specific domains.

Math. Two recent works addressed how formulas conveyed in text should be accommodated within AMR graphs. Iordan (2021) developed an AMR parser with an added coreference detection feature to parse AMR graphs from descriptions of geometry problems, and Mansouri et al. (2022) incorporated embeddings of altered AMRs into an LLM for extracting formulas.

Legal Reasoning. As is the case for math, domain-specific language in legal documents necessitates alteration of the AMR formalism. Vu and Nguyen (2019) evaluated the (generally poor) performance of AMR parsers on legal documents; further, Schrack et al. (2022) showed that neuro-symbolic methods which include linearized AMR graphs do not outperform text-only methods on multiple choice question answering for legal reasoning, but do offer a complementary signal. To address these challenges, Vu et al. (2022) introduced a human-annotated dataset of AMRs in the legal domain.

Spatial/Situated Dialogue. A fruitful line of AMR application research has focused on spatial/ situated dialogue, in particular on **human-robot interaction**.[3] Numerous datasets of AMR graphs of human-robot interactions have been created (Bastianelli et al., 2014; Shichman et al., 2023). An altered AMR schema called "Dialogue-AMR" (containing information on tense, aspect, and speech acts) supports the representation of human-robot interactions (Bonial et al., 2019; Abrams et al., 2020; Bonial et al., 2020a, 2021, 2023). Ultimately, this has enabled grounded natural language understanding for human-robot interactions.

Other work in spatial and situated AMR (unrelated to human-robot interaction) has also accounted for the necessity of altering AMR to include **grounding language**. Datasets of AMR graphs for multimodal dialogue have incorporated gestures (Donatelli et al., 2022; Lai et al., 2024) and spatial information (Bonn et al., 2020; Dan et al., 2020) into the AMR schema. Martin et al.

[3]While AMR-based dialogue understanding work has been primarily focused on human-robot dialogue, Bai et al. (2022) achieved state-of-the-art performance on general dialogue understanding by using AMR to continuously pre-train a Transformer encoder.

(2020) crowdsourced AMR annotations of text containing spatial information and Tam et al. (2023) investigated action annotation in AMR.

System Requirements. Lamercerie and Foret (2021) altered AMR graphs of system requirements by grouping subgraphs of individual desired properties of the system.

Recipe Instructions. Similarly to the work done on system requirements, Stein et al. (2023) modified AMR graphs for recipes, breaking down sentence-level AMR graphs into graphs of individual actions in the recipe.

3.4 AMR for Image and Speech

As a semantic representation, AMR has been converted into other types of text-based formalisms (such as SPARQL in §3.1 and UMR (Post et al., 2024)), as well as leveraged in support of non-text-based forms of media such as images and speech.

Images. Recent work has investigated the use of AMR for **scene graph parsing**, which is the production of a graph-based representation of object boundaries in images. Choi et al. (2022a,b) converted AMR graphs into scene graphs, while Abdelsalam et al. (2022) explored the use of AMR *as an alternative* to scene graphs (via image-to-AMR parsing).

Image captioning has employed AMR in order to focus on specific aspects of meaning or task-specific difficulties. Neto et al. (2020) used AMR to produce descriptions of specific regions of an image, and Kim et al. (2024) used the relationship between the sentence and object (via AMR) for caption debiasing. Finally, Bhattacharyya et al. (2024) and Chen et al. (2024) leveraged semantic relations from AMR and the image to guide caption generation.

Speech. Little work has investigated the utility of AMR for speech systems, though Addlesee and Damonte (2023) addressed nonstandard speech as an accessibility issue for voice assistants by producing a corpus of AMR graphs of disrupted speech, and training models on this data.

3.5 AMR for NLG Evaluation

AMR's nature as an interpretable semantic representation lends itself to evaluation-based tasks.

Dialogue Evaluation. Ghazarian et al. (2022) developed a robust dialogue coherence measure by training on negative text examples that are generated from AMRs which were manipulated in controlled ways, e.g., introducing contradictions by changing node labels to antonyms. The resulting model achieves significantly better correlation to humans than other text- and graph-based baselines. On the same task, Yang et al. (2024) showed that performance improvements can also be achieved by fusing text and AMR in a model using a dual-encoder, thus more directly using the AMR.

Summary Evaluation. Ribeiro et al. (2022) trained a model that leverages AMR as auxiliary information in a dual encoder, outperforming strong QA-based and NLI-based **summary factuality** models. Addressing the same task, Qiu et al. (2024) produced training examples with manipulated AMRs, resulting in a state-of-the-art factuality prediction model. Tackling the second pillar of summary quality, being **summary relevance**, Nawrath et al. (2024) split AMR graphs of summaries into subgraphs with the aim to generate Summary Content Units (clauses that identify sub-sentential content in summaries (Nenkova and Passonneau, 2004)). In this case, the results were more mixed and the authors noted that development of advanced splitting methods is necessary for improved results. As a general **summary interpretability** method, Landes and Di Eugenio (2024) developed an AMR-alignment tool for the inspection of summaries, aligning parts of the summary with the evidence in the source document.

General Evaluation and Diagnostics. Opitz and Frank (2021) used AMR metrics to compare AMR graphs of candidates and references, enabling measurement of fine-grained text quality aspects like polarity or coreference faithfulness. Using AMR metrics to evaluate NLG quality is limited by current parsing inaccuracies (Manning and Schneider, 2021).

3.6 AMR for Language Studies

AMR is a linguistic tool which has been utilized for language-focused research and teaching.

Linguistic Research. Sawai et al. (2015) used AMR to build a model that answers statistical research questions about the semantic structure of noun phrases.

Teaching. The investigation of the meaning of a text emerges as an intriguing and interesting classroom exercise. In particular, given its linguistic specificity and interpretability, AMR can help students learn about linguistic structures, as exemplified in the lesson and exercise on AMR in Eisenstein (2021).

3.7 AMR for Explainable Semantic Similarity

AMR-based metrics are of wider interest in measuring semantic similarity and relatedness, beyond NLG evaluation (c.f. §3.5). Intuitively, we can parse two input texts and calculate AMR similarity, providing an additional layer of interpretability and explainability via AMR.

AMR metrics have been used for detecting paraphrases (Issa et al., 2018), evaluating the answers provided by language learners on reading comprehension questions (Dellert, 2020), judging argument and text similarity (Opitz et al., 2021b), and matching local knowledge graphs (Kachwala et al., 2024). Furthermore, assessing structural graph isomorphism in the AMRs of multilingual texts achieves finer-grained semantic equivalence judgments than neural methods (Wein et al., 2023). Incorporating AMR graphs and AMR metrics into neural models for **natural language inference** has also been of value (Opitz et al., 2023; Feng and Hunter, 2024; Bao et al., 2023).

Neural text embedding models such as SBERT (Reimers and Gurevych, 2019) and SimCSE (Gao et al., 2021) have been retro-fitted by encoding AMR graphs (Cai et al., 2022). Alternatively, semantic **embedding interpretability** has been induced by binding parts of embeddings to semantic features such as negation, semantic roles, or named entities, that can be measured with AMR metrics (Opitz and Frank, 2022b). In the same direction, Fodor et al. (2024) found that state-of-the-art transformers "poorly capture the pattern of human semantic similarity judgments", and AMR can be used to build simple methods that combine semantic compontents into an improved hybrid model.

3.8 Miscellanea

Finally, we discuss miscellanea, which are either applications where the impetus behind the use of AMR may be less obvious, or applications that escape a categorization into the above classes. Firstly, AMRs have been employed for **commonsense reasoning**, using different strategies: tracing reasoning paths through AMRs (Lim et al., 2020), enriching AMRs with relations from a Commonsense Knowledge Graph (Oh et al., 2022), and within a neuro-symbolic approach where AMR is converted into first-order logic (Chanin and Hunter, 2023). AMR has also been used for **sentiment analysis** (Ma et al., 2023) and to generate feedback for **reinforcement learning** in text-based games (Chaudhury et al., 2023). Elbasani and Kim (2022) parsed AMRs of the text and then used that as input to a convolutional neural network for **toxic content detection**. For a similar task–**fake news detection**–Gupta et al. (2023) used text-based features in conjunction with AMR graphs to classify whether a tweet is fake news. Finally, AMR has been used to perform general **text classification** (Ogawa and Saga, 2023).

4 Engineering with AMR

In the prior section, we categorized AMR applications by task in order to provide an overview of the AMR application landscape, showcasing AMR as a general-purpose representation. In this section, we describe and provide a functional guide to the techniques and patterns which have allowed AMR to be leveraged for the aforementioned engineering purposes.

4.1 AMR Preparation

As an initial step in working with AMR, many applications conduct operations on the AMR graph. We observe frequent use of the following three types of operations: pre-processing, splitting/merging, and encoding.

AMR pre-processing can range from simple string changes to more elaborate graph transformations. Examples of simple string changes include lower-casing or truncating the concept labels. **Graph transformations** that preserve the equivalency of AMRs can include *reification* (Opitz et al., 2021a; Shou and Lin, 2023), where, with the help of a dictionary, we 'generalize' binary edge labels to n-ary structures. Alternatively conversion to a *Levi Graph* (Beck et al., 2018; Lim et al., 2020) which is a bipartite graph without edge labels, alleviates the need to handle edge labels in some specific way other than node labels (see Appendix A for examples of these transformations).

AMR splitting and merging can also come in handy. For example, AMRs are split to find the largest common sub-structures in question answer pairs (Deng et al., 2022), or to extract subgraphs that elicit specific aspects of meaning such as polarity or semantic roles (Opitz and Frank, 2022b; Opitz, 2023). Merging can be applied by first matching concepts or named entities from two graphs, and then connecting or fusing nodes that represent the same entities (Liu et al., 2015; Lee et al., 2021), possibly leveraging advanced coreference resolution within AMR (Fu et al., 2021). In the simplest case, merging is conducted by connecting multiple graphs at their roots (Kouris et al., 2022; Bai et al., 2022).

Bai et al. (2022) also exemplifies the possibility of **AMR enrichment** with task-specific information (here: edges labeled with the speaker in a dialogue). Other examples of additional information used to enrich AMR graphs include VerbNet event structure (Tu et al., 2024) and links from knowledge graphs (Zhang et al., 2021).

These line of work on AMR merging and enrichment may profit from the ongoing research into the 'AMR-intrinsic' tasks of AMR coreference resolution (Fu et al., 2021; Li et al., 2022) and AMR-to-text alignment (Blodgett and Schneider, 2021; Martínez Lorenzo et al., 2023).

Many approaches have required that AMR somehow be encoded into an external model. Synergizing well with the strong NLU inductive bias of text language models, one successful paradigm for **AMR encoding** is to simply feed the linearized graph as a string, where string pre-processing tricks (such as those described) can increase performance (Ribeiro et al., 2021b,a).

AMR encoding can also involve constructing feature vectors (/embeddings) of the full AMR graphs (Wang et al., 2017) or targeted semantic parts (Fodor et al., 2024).

Prior work on AMR-to-text generation has found success encoding AMRs using Graph RNNs (Song et al., 2018) and Graph Transformers (Song et al., 2020; Yao et al., 2020), and the same or similar encoding mechanisms are also found when encoding AMRs for downstream applications (Song et al., 2019).

4.2 Two Processing Paradigms

We observe two major processing paradigms in AMR applications: the neuro-symbolic model, and the use of AMR as an intermediate representation. The first approach has been consistently popular; the second approach has grown in popularity more recently.

Fusing Text and AMR: the Neuro-symbolic Model. The abundant recent interest in neuro-symbolic approaches for NLP (Besold et al., 2021; Hamilton et al., 2022; Yu et al., 2023) has bled into AMR applications.

A common way of leveraging AMR information is merging information from the AMR modality with information from the text modality, typically with an auxiliary motive (e.g., AMR is used to help refine the extracted information from the text to improve a model accuracy by some points). To accomplish this, a prominent strategy has been to construct an AMR parse from the text and then feed both this parse and text into one neural model.

Sometimes, a joint encoder is employed, where AMR and text are simply concatenated and fused at the lowest processing layers (Huang et al., 2022; Hsu et al., 2023). The two modalities (text and AMR) can also be first processed separately, using two individual encoders, to create disjoint higher-level representations that are then fused later such as by adding or concatenating. This fusing can happen in intermediate layers (Dai et al., 2022; Ma et al., 2023), or at the final decision layer (Cai et al., 2022; Opitz et al., 2023).

AMR as an Intermediate Representation. Using AMR as an intermediate representation means typically operating on and with the AMR X as follows: **parse** $\rightarrow X \rightarrow$ **generate**, interlinking parsing and generation models.

One appealing aspect of this technique is the increased interpretility and linguistic control, as to **induce controlled changes in meaning**. For example, the AMR graph can be transformed to generate (1) paraphrases (e.g., by swapping the root (Huang et al., 2023), or swapping out a concept with a synonym, and then generating text (Shi et al., 2023)), or to generate (2) contradictions (e.g., by inserting a targeted negation to a predicate and then generating text from the manipulated structure (Ghazarian et al., 2022)).

On the other hand, the AMR graph can instead remain unaltered while the input text, parsing method, or generation method are varied, such as in the cases of Jangra et al. (2022) for style transfer and Wein and Schneider (2024b) for translationese reduction. As another example, Dohare et al. (2017) compile a summary AMR, by finding AMR nodes focused on important entities, and selecting the subtree hanging from that verb as the summary AMR. Text is then generated from the specified subtree. This highlights that the splitting and merging techniques highlighted in §4.1 can be part of working with AMR as intermediate representation.

5 Areas for Future Work

Surveying the vast number of tasks and techniques utilized throughout the last decade, we observe three notable areas for future work on AMR applications.

First, a technique which has shown great promise for incorporating AMR into neural or non-neural downstream applications is as an **intermediate representation**. This intuitively leverages both AMR's design as a graph-based semantic representation as well as the progress on text-to-AMR parsing and AMR-to-text generation.[4] Using AMR as an intermediary provides us linguistic control and interpretability, which are increasingly desirable in the age of "black box" neural models. Numerous recent studies have successfully exploited AMR as an intermediary (§4.2), indicating that this may be a promising path forward, particularly in low-resource settings or for data augmentation.

Second, recent work has shown the benefits of incorporating AMR in **few-shot or low-resource settings** (Nguyen et al., 2021; Gururaja et al., 2023; Hua et al., 2023; Ghosh et al., 2024). This indicates that, regardless of the technique of incorporation, AMR is positioned to be especially well suited for engineering gains in these settings.

Finally, an area which has been largely understudied (§3.6) but directly follows from the design of AMR as semantic representation, is the use of AMR for **linguistic analysis** and text statistics. Potential applications include language learning or studying predicate-argument patterns in L1 or L2 texts. Additional uses of AMR for linguistic analysis include linguistically-focused evaluations and finer statistics of NLG systems (prior work in this direction discussed in §3.5).[5]

6 Related Work on Applications of Other Meaning Representations

Other surveys have considered the engineering utility of semantic representations. Regarding specific tasks, Verrev (2023) tested out the benefits of various meaning representations for knowledge-base question answering (KBQA), and Prange et al. (2022) for next-word prediction in conjunction with neural models.

Related work has also compared the designs and features of semantic representations (Abend and Rappoport, 2017), with Pavlova et al. (2023) specifically addressing how the features of the semantic representations may play a role in their utility.

7 Conclusion

In this survey, we provided a thorough overview of the tasks where Abstract Meaning Representation graphs have been used and the techniques involved in using AMR for engineering purposes. Given the availability of strong parsing systems and the increased interest in AMR, we expect that we are on the precipice of exciting progress employing AMR in downstream applications. As our synthesis of AMR engineering patterns indicates, there are numerous methods, techniques and possible applications that await further exploration and continued improvement.

Limitations

In this survey, we direct our attention exclusively towards the AMR formalism given the recent abundance of work incorporating AMR, and the fact that there are still few surveys addressing AMR. While other semantic representations have been considered as engineering tools (§6), AMR is currently unique in the breadth with which it has been used and studied.

As discussed in §2.1, applications of AMR have been primarily focused on English; recent work (discussed throughout this survey) has demonstrated the cross-lingual and non-English utility of AMR (Wein and Schneider, 2024a), which continues to increase given advancements in multilingual AMR parsing and generation.

We have incorporated the full breadth of existing AMR application work to the best of our knowledge; the bar chart in §3 serves as a lower bound as there may be papers that were missed.

Acknowledgements

We thank anonymous reviewers for their helpful feedback and recommendations.

References

Mohamed Ashraf Abdelsalam, Zhan Shi, Federico Fancellu, Kalliopi Basioti, Dhaivat Bhatt, Vladimir Pavlovic, and Afsaneh Fazly. 2022. Visual semantic parsing: From images to Abstract Meaning Representation. In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 282–300, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

[4]This approach also hearkens back to one of the classic approaches towards a fundamental NLP problem, being machine translation: interlingual machine translation (Dorr, 1993).

[5]For mining large unstructured text data, AMR offers semantic triplets that await to be sensibly aggregated, e.g., to craft an automatic knowledge graph.

Omri Abend and Ari Rappoport. 2017. The state of the art in semantic representation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 77–89, Vancouver, Canada. Association for Computational Linguistics.

Mitchell Abrams, Claire Bonial, and Lucia Donatelli. 2020. Graph-to-graph meaning representation transformations for human-robot dialogue. In *Proceedings of the Society for Computation in Linguistics 2020*, pages 250–253, New York, New York. Association for Computational Linguistics.

Angus Addlesee and Marco Damonte. 2023. Understanding disrupted sentences using underspecified abstract meaning representation. In *Interspeech 2023*.

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).

Zahra Azin and Gülşen Eryiğit. 2019. Towards Turkish Abstract Meaning Representation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 43–47, Florence, Italy. Association for Computational Linguistics.

Xuefeng Bai, Linfeng Song, and Yue Zhang. 2022. Semantic-based pre-training for dialogue understanding. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 592–607, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Mollie Bakal. 2021. Graph-to-graph translations to augment abstract meaning representation tense and aspect. Bachelor's thesis, University of Michigan.

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.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2019. Abstract Meaning Representation (AMR) 1.2.6 specification. `https://github.com/amrisi/amr-guidelines/blob/master/amr.md`.

Qiming Bao, Alex Yuxuan Peng, Zhenyun Deng, Wanjun Zhong, Neset Tan, Nathan Young, Yang Chen, Yonghua Zhu, Michael Witbrock, and Jiamou Liu. 2023. Contrastive learning with logic-driven data augmentation for logical reasoning over text. *arXiv preprint arXiv:2305.12599*.

Jorge Baptista, Sónia Reis, João Dias, and Pedro Santos. 2024. Lexicalized meaning representation (LMR). In *Proceedings of the Fifth International Workshop on Designing Meaning Representations @ LREC-COLING 2024*, pages 101–111, Torino, Italia. ELRA and ICCL.

Emanuele Bastianelli, Giuseppe Castellucci, Danilo Croce, Luca Iocchi, Roberto Basili, and Daniele Nardi. 2014. HuRIC: a human robot interaction corpus. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4519–4526, Reykjavik, Iceland. European Language Resources Association (ELRA).

Daniel Beck, Gholamreza Haffari, and Trevor Cohn. 2018. Graph-to-sequence learning using gated graph neural networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 273–283, Melbourne, Australia. Association for Computational Linguistics.

Tarek R Besold, Artur d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kühnberger, Luis C Lamb, Priscila Machado Vieira Lima, Leo de Penning, et al. 2021. Neural-symbolic learning and reasoning: A survey and interpretation 1. In *Neuro-Symbolic Artificial Intelligence: The State of the Art*, pages 1–51. IOS press.

Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021. One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12564–12573.

Abhidip Bhattacharyya, Martha Palmer, and Christoffer Heckman. 2024. ReCAP: Semantic role enhanced caption generation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13633–13649, Torino, Italy. ELRA and ICCL.

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.

Claire Bonial, Mitchell Abrams, David Traum, and Clare Voss. 2021. Builder, we have done it: Evaluating & extending dialogue-AMR NLU pipeline for two collaborative domains. In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*, pages 173–183, Groningen, The Netherlands (online). Association for Computational Linguistics.

Claire Bonial, Lucia Donatelli, Mitchell Abrams, Stephanie M. Lukin, Stephen Tratz, Matthew Marge, Ron Artstein, David Traum, and Clare Voss. 2020a. Dialogue-AMR: Abstract Meaning Representation for dialogue. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 684–695, Marseille, France. European Language Resources Association.

Claire Bonial, Lucia Donatelli, Stephanie M. Lukin, Stephen Tratz, Ron Artstein, David Traum, and Clare Voss. 2019. Augmenting Abstract Meaning Representation for human-robot dialogue. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 199–210, Florence, Italy. Association for Computational Linguistics.

Claire Bonial, Julie Foresta, Nicholas C. Fung, Cory J. Hayes, Philip Osteen, Jacob Arkin, Benned Hedegaard, and Thomas Howard. 2023. Abstract Meaning Representation for grounded human-robot communication. In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 34–44, Nancy, France. Association for Computational Linguistics.

Claire Bonial, Stephanie M. Lukin, David Doughty, Steven Hill, and Clare Voss. 2020b. InfoForager: Leveraging semantic search with AMR for COVID-19 research. In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 67–77, Barcelona Spain (online). Association for Computational Linguistics.

Julia Bonn, Martha Palmer, Zheng Cai, and Kristin Wright-Bettner. 2020. Spatial AMR: Expanded spatial annotation in the context of a grounded Minecraft corpus. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4883–4892, Marseille, France. European Language Resources Association.

Mihaela Bornea, Ramon Fernandez Astudillo, Tahira Naseem, Nandana Mihindukulasooriya, Ibrahim Abdelaziz, Pavan Kapanipathi, Radu Florian, and Salim Roukos. 2021. Learning to transpile AMR into SPARQL. *arXiv preprint arXiv:2112.07877*.

Gully A. Burns, Ulf Hermjakob, and José Luis Ambite. 2016. Abstract meaning representations as linked data. In *The Semantic Web – ISWC 2016*, pages 12–20. Springer International Publishing.

Deng Cai, Xin Li, Jackie Chun-Sing Ho, Lidong Bing, and Wai Lam. 2022. Retrofitting multilingual sentence embeddings with Abstract Meaning Representation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6456–6472, Abu Dhabi, United Arab Emirates. 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.

Hejing Cao and Dongyan Zhao. 2023. Leveraging denoised Abstract Meaning Representation for grammatical error correction. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7180–7188, Toronto, Canada. Association for Computational Linguistics.

David Chanin and Anthony Hunter. 2023. Neuro-symbolic commonsense social reasoning. *arXiv preprint arXiv:2303.08264*.

Subhajit Chaudhury, Sarathkrishna Swaminathan, Daiki Kimura, Prithviraj Sen, Keerthiram Murugesan, Rosario Uceda-Sosa, Michiaki Tatsubori, Achille Fokoue, Pavan Kapanipathi, Asim Munawar, and Alexander Gray. 2023. Learning symbolic rules over Abstract Meaning Representations for textual reinforcement learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6764–6776, Toronto, Canada. Association for Computational Linguistics.

Feng Chen, Xinyi Li, Jintao Tang, Shasha Li, and Ting Wang. 2024. Benefit from AMR: Image captioning withexplicit relations and endogenous knowledge. In *Web and Big Data*, pages 363–376, Singapore. Springer Nature Singapore.

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.

Hyonsu Choe, Jiyoon Han, Hyejin Park, Tae Hwan Oh, and Hansaem Kim. 2020. Building Korean Abstract Meaning Representation corpus. In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 21–29, Barcelona Spain (online). Association for Computational Linguistics.

Woo Suk Choi, Yu-Jung Heo, Dharani Punithan, and Byoung-Tak Zhang. 2022a. Scene graph parsing via Abstract Meaning Representation in pre-trained language models. In *Proceedings of the 2nd Workshop on Deep Learning on Graphs for Natural Language Processing (DLG4NLP 2022)*, pages 30–35, Seattle, Washington. Association for Computational Linguistics.

Woo Suk Choi, Yu-Jung Heo, and Byoung-Tak Zhang. 2022b. Sgram: Improving scene graph parsing via abstract meaning representation. *arXiv preprint arXiv:2210.08675*.

Xinbang Dai, Huiying Li, Tenggou Wang, Lei Li, Yuxi Feng, and Xin Li. 2022. Sparel: A semantic parsing relation linking method for knowledge base question answering. In *Proceedings of the 11th International Joint Conference on Knowledge Graphs*, pages 73–81.

Soham Dan, Parisa Kordjamshidi, Julia Bonn, Archna Bhatia, Zheng Cai, Martha Palmer, and Dan Roth. 2020. From spatial relations to spatial configurations. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5855–5864, Marseille, France. European Language Resources Association.

Johannes Dellert. 2020. Exploring probabilistic soft logic as a framework for integrating top-down and bottom-up processing of language in a task context. *arXiv preprint arXiv:2004.07000*.

Zhenyun Deng, Yonghua Zhu, Yang Chen, Michael Witbrock, and Patricia Riddle. 2022. Interpretable AMR-based question decomposition for multi-hop question answering. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 4093–4099. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Shibhansh Dohare, Harish Karnick, and Vivek Gupta. 2017. Text summarization using abstract meaning representation. *arXiv preprint arXiv:1706.01678*.

Lucia Donatelli, Kenneth Lai, Richard Brutti, and James Pustejovsky. 2022. Towards situated AMR: Creating a corpus of gesture AMR. In *International Conference on Human-Computer Interaction*, pages 293–312. Springer.

Lucia Donatelli, Michael Regan, William Croft, and Nathan Schneider. 2018. Annotation of tense and aspect semantics for sentential AMR. In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 96–108, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Bonnie J Dorr. 1993. Interlingual machine translation: A parameterized approach. *Artificial Intelligence*, 63:429 – 492.

Jacob Eisenstein. 2021. On writing a textbook on natural language processing. In *Proceedings of the Fifth Workshop on Teaching NLP*, pages 125–130, Online. Association for Computational Linguistics.

Ermal Elbasani and Jeong-Dong Kim. 2022. Amr-cnn: Abstract meaning representation with convolution neural network for toxic content detection. *Journal of Web Engineering*, 21(03):677–692.

Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. The MIT Press.

Xuyao Feng and Anthony Hunter. 2024. Identification of entailment and contradiction relations between natural language sentences: A neurosymbolic approach. *arXiv preprint arXiv:2405.01259*.

James Fodor, Simon De Deyne, and Shinsuke Suzuki. 2024. Compositionality and Sentence Meaning: Comparing Semantic Parsing and Transformers on a Challenging Sentence Similarity Dataset. *Computational Linguistics*, pages 1–52.

Giacomo Frisoni, Paolo Italiani, Stefano Salvatori, and Gianluca Moro. 2023. Cogito ergo summ: Abstractive summarization of biomedical papers via semantic parsing graphs and consistency rewards. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):12781–12789.

Qiankun Fu, Linfeng Song, Wenyu Du, and Yue Zhang. 2021. End-to-end AMR coreference resolution. 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 4204–4214, Online. Association for Computational Linguistics.

Boris Galitsky. 2020. *Employing Abstract Meaning Representation to Lay the Last-Mile Toward Reading Comprehension*, pages 57–86. Springer International Publishing, Cham.

Aldo Gangemi, Arianna Graciotti, Antonello Meloni, Andrea Nuzzolese, Valentina Presutti, Diego Reforgiato Recupero, Alessandro Russo, and Rocco Tripodi. 2023. Text2AMR2FRED, a tool for transforming text into rdf/owl knowledge graphs via abstract meaning representation.

Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Sahil Garg, Aram Galstyan, Ulf Hermjakob, and Daniel Marcu. 2016. Extracting biomolecular interactions using semantic parsing of biomedical text. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.

Sarik Ghazarian, Nuan Wen, Aram Galstyan, and Nanyun Peng. 2022. DEAM: Dialogue coherence evaluation using AMR-based semantic manipulations. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 771–785, Dublin, Ireland. Association for Computational Linguistics.

Sreyan Ghosh, Utkarsh Tyagi, Sonal Kumar, Chandra Kiran Evuru, Ramaneswaran S, S Sakshi, and Dinesh Manocha. 2024. ABEX: Data augmentation for low-resource NLU via expanding abstract descriptions. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 726–748, Bangkok, Thailand. Association for Computational Linguistics.

Jonas Groschwitz, Shay Cohen, Lucia Donatelli, and Meaghan Fowlie. 2023. AMR parsing is far from solved: GrAPES, the granular AMR parsing evaluation suite. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10728–10752, Singapore. Association for Computational Linguistics.

Shubham Gupta, Narendra Yadav, Suman Kundu, and Sainathreddy Sankepally. 2023. FakEDAMR: Fake news detection using abstract meaning representation network. In *International Conference on Complex Networks and Their Applications*, pages 308–319. Springer.

Sireesh Gururaja, Ritam Dutt, Tinglong Liao, and Carolyn Rosé. 2023. Linguistic representations for fewer-shot relation extraction across domains. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7502–7514, Toronto, Canada. Association for Computational Linguistics.

Kyle Hamilton, Aparna Nayak, Bojan Božić, and Luca Longo. 2022. Is neuro-symbolic AI meeting its promises in natural language processing? a structured review. *Semantic Web*, pages 1–42.

Hardy Hardy and Andreas Vlachos. 2018. Guided neural language generation for abstractive summarization using Abstract Meaning Representation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 768–773, Brussels, Belgium. Association for Computational Linguistics.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

I-Hung Hsu, Zhiyu Xie, Kuan-Hao Huang, Prem Natarajan, and Nanyun Peng. 2023. AMPERE: AMR-aware prefix for generation-based event argument extraction model. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10976–10993, Toronto, Canada. Association for Computational Linguistics.

Yilun Hua, Zhaoyuan Deng, and Kathleen McKeown. 2023. Improving long dialogue summarization with semantic graph representation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13851–13883, Toronto, Canada. Association for Computational Linguistics.

Yilun Hua, Zhaoyuan Deng, and Zhijie Xu. 2022. AMRTVSumm: AMR-augmented hierarchical network for TV transcript summarization. In *Proceedings of The Workshop on Automatic Summarization for Creative Writing*, pages 36–43, Gyeongju, Republic of Korea. Association for Computational Linguistics.

Kuan-Hao Huang, Varun Iyer, I-Hung Hsu, Anoop Kumar, Kai-Wei Chang, and Aram Galstyan. 2023. ParaAMR: A large-scale syntactically diverse paraphrase dataset by AMR back-translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8047–8061, Toronto, Canada. Association for Computational Linguistics.

Kuan-Hao Huang, Varun Iyer, Anoop Kumar, Sriram Venkatapathy, Kai-Wei Chang, and Aram Galstyan. 2022. Unsupervised syntactically controlled paraphrase generation with Abstract Meaning Representations. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1547–1554, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Marcio Inácio and Thiago Pardo. 2021. Semantic-based opinion summarization. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 619–628, Held Online. INCOMA Ltd.

Marcio Lima Inácio, Marco Antonio Sobrevilla Cabezudo, Renata Ramisch, Ariani Di Felippo, and Thiago Alexandre Salgueiro Pardo. 2022. The AMR-PT corpus and the semantic annotation of challenging sentences from journalistic and opinion texts. *SciELO Preprints*.

Anca-Elena Iordan. 2021. Automatic comprehension of geometry problems using AMR parser. In *International Conference on Software Engineering and Knowledge Engineering (SEKE)*, pages 628–631.

Md Saiful Islam, Adiba Proma, Yilin Zhou, Syeda Nahida Akter, Caleb Wohn, and Ehsan Hoque. 2022. KnowUREnvironment: An automated knowledge graph for climate change and environmental issues. In *AAAI 2022 Fall Symposium: The Role of AI in Responding to Climate Challenges*.

Fuad Issa, Marco Damonte, Shay B. Cohen, Xiaohui Yan, and Yi Chang. 2018. Abstract Meaning Representation for paraphrase detection. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 442–452, New Orleans, Louisiana. Association for Computational Linguistics.

Anubhav Jangra, Preksha Nema, and Aravindan Raghuveer. 2022. T-STAR: Truthful style transfer using AMR graph as intermediate representation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8805–8825, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Zhijing Jin, Yuen Chen, Fernando Gonzalez Adauto, Jiarui Liu, Jiayi Zhang, Julian Michael, Bernhard Schölkopf, and Mona Diab. 2024. Analyzing the role of semantic representations in the era of large language models. 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

3781–3798, Mexico City, Mexico. Association for Computational Linguistics.

Zoher Kachwala, Jisun An, Haewoon Kwak, and Filippo Menczer. 2024. REMATCH: Robust and efficient matching of local knowledge graphs to improve structural and semantic similarity. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1018–1028, Mexico City, Mexico. Association for Computational Linguistics.

Pavan Kapanipathi, Ibrahim Abdelaziz, Srinivas Ravishankar, Salim Roukos, Alexander Gray, Ramón Fernandez Astudillo, Maria Chang, Cristina Cornelio, Saswati Dana, Achille Fokoue, Dinesh Garg, Alfio Gliozzo, Sairam Gurajada, Hima Karanam, Naweed Khan, Dinesh Khandelwal, Young-Suk Lee, Yunyao Li, Francois Luus, Ndivhuwo Makondo, Nandana Mihindukulasooriya, Tahira Naseem, Sumit Neelam, Lucian Popa, Revanth Gangi Reddy, Ryan Riegel, Gaetano Rossiello, Udit Sharma, G P Shrivatsa Bhargav, and Mo Yu. 2021. Leveraging Abstract Meaning Representation for knowledge base question answering. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3884–3894, Online. Association for Computational Linguistics.

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*.

Jungeun Kim, Jangyeong Jeon, Seungjin Jung, and Junyeong Kim. 2024. De-bias using abstract meaning representation for image captioning. In *2024 IEEE International Conference on Consumer Electronics (ICCE)*, pages 1–4. IEEE.

Panagiotis Kouris, Georgios Alexandridis, and Andreas Stafylopatis. 2022. Text summarization based on semantic graphs: An abstract meaning representation graph-to-text deep learning approach.

Kenneth Lai, Richard Brutti, Lucia Donatelli, and James Pustejovsky. 2024. Encoding gesture in multimodal dialogue: Creating a corpus of multimodal AMR. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5806–5818, Torino, Italy. ELRA and ICCL.

Aurélien Lamercerie and Annie Foret. 2021. From requirements as AMR-like graphs to automata-based reasoning. Research report, Université de Rennes 1.

Paul Landes and Barbara Di Eugenio. 2024. CALAMR: Component ALignment for Abstract Meaning Representation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2622–2637, Torino, Italia. ELRA and ICCL.

Fei-Tzin Lee, Miguel Ballesteros, Feng Nan, and Kathleen McKeown. 2022a. Using structured content plans for fine-grained syntactic control in pretrained language model generation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5882–5895, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Fei-Tzin Lee, Chris Kedzie, Nakul Verma, and Kathleen McKeown. 2021. An analysis of document graph construction methods for AMR summarization. *arXiv preprint arXiv:2111.13993*.

Young-Suk Lee, Ramón Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian, and Salim Roukos. 2022b. 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.

Bin Li, Yuan Wen, Weiguang Qu, Lijun Bu, and Nianwen Xue. 2016. Annotating the little prince with Chinese AMRs. In *Proceedings of the 10th Linguistic Annotation Workshop held in conjunction with ACL 2016 (LAW-X 2016)*, pages 7–15, Berlin, Germany. Association for Computational Linguistics.

Changmao Li and Jeffrey Flanigan. 2022. Improving neural machine translation with the Abstract Meaning Representation by combining graph and sequence transformers. In *Proceedings of the 2nd Workshop on Deep Learning on Graphs for Natural Language Processing (DLG4NLP 2022)*, pages 12–21, Seattle, Washington. Association for Computational Linguistics.

Irene Li, Linfeng Song, Kun Xu, and Dong Yu. 2022. Variational graph autoencoding as cheap supervision for AMR coreference resolution. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2790–2800, Dublin, Ireland. Association for Computational Linguistics.

Xiang Li, Thien Huu Nguyen, Kai Cao, and Ralph Grishman. 2015. Improving event detection with Abstract Meaning Representation. In *Proceedings of the First Workshop on Computing News Storylines*, pages 11–15, Beijing, China. Association for Computational Linguistics.

Kexin Liao, Logan Lebanoff, and Fei Liu. 2018. Abstract Meaning Representation for multi-document summarization. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1178–1190, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Jungwoo Lim, Dongsuk Oh, Yoonna Jang, Kisu Yang, and Heuiseok Lim. 2020. I know what you asked: Graph path learning using AMR for commonsense reasoning. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2459–2471, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Ha Linh and Huyen Nguyen. 2019. A case study on meaning representation for Vietnamese. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 148–153, Florence, Italy. Association for Computational Linguistics.

Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. Toward abstractive summarization using semantic representations. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1077–1086, Denver, Colorado. Association for Computational Linguistics.

Fukun Ma, Xuming Hu, Aiwei Liu, Yawen Yang, Shuang Li, Philip S. Yu, and Lijie Wen. 2023. AMR-based network for aspect-based sentiment analysis. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 322–337, Toronto, Canada. Association for Computational Linguistics.

Emma Manning and Nathan Schneider. 2021. Referenceless parsing-based evaluation of AMR-to-English generation. In *Proceedings of the 2nd Workshop on Evaluation and Comparison of NLP Systems*, pages 114–122, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Behrooz Mansouri, Ricardo Campos, and Adam Jatowt. 2023. Towards timeline generation with abstract meaning representation. In *Companion Proceedings of the ACM Web Conference 2023*, WWW '23 Companion, page 1204–1207, New York, NY, USA. Association for Computing Machinery.

Behrooz Mansouri, Douglas W. Oard, and Richard Zanibbi. 2022. Contextualized formula search using math abstract meaning representation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, CIKM '22, page 4329–4333, New York, NY, USA. Association for Computing Machinery.

Mary Martin, Cecilia Mauceri, Martha Palmer, and Christoffer Heckman. 2020. Leveraging non-specialists for accurate and time efficient AMR annotation. In *Proceedings of the LREC 2020 Workshop on "Citizen Linguistics in Language Resource Development"*, pages 35–39, Marseille, France. European Language Resources Association.

Abelardo Carlos Martínez Lorenzo, Pere Lluís Huguet Cabot, and Roberto Navigli. 2023. Cross-lingual AMR aligner: Paying attention to cross-attention. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1726–1742, Toronto, Canada. Association for Computational Linguistics.

Abelardo Carlos Martínez Lorenzo, Marco Maru, and Roberto Navigli. 2022. Fully-Semantic Parsing and Generation: the BabelNet Meaning Representation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1727–1741, Dublin, Ireland. Association for Computational Linguistics.

Antonello Meloni, Diego Reforgiato Recupero, and Aldo Gangemi. 2017. AMR2FRED, a tool for translating abstract meaning representation to motif-based linguistic knowledge graphs. In *The Semantic Web: ESWC 2017 Satellite Events*, pages 43–47, Cham. Springer International Publishing.

Noelia Migueles-Abraira, Rodrigo Agerri, and Arantza Diaz de Ilarraza. 2018. Annotating Abstract Meaning Representations for Spanish. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

George A. Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

Tuosiyu Ming, Hongchang Chen, and Yizhuo Yang. 2018. A method of semantic redundant information filtering for abstract meaning representation graph. In *2018 2nd International Conference on Data Science and Business Analytics (ICDSBA)*, pages 367–372. IEEE.

Arindam Mitra and Chitta Baral. 2016. Addressing a question answering challenge by combining statistical methods with inductive rule learning and reasoning. volume 30.

Paloma Moreda, Armando Suárez, Elena Lloret, Estela Saquete, and Isabel Moreno. 2018. From sentences to documents: Extending abstract meaning representation for understanding documents. *Procesamiento del Lenguaje Natural*, 60:61–68.

Almuth Müller and Achim Kuwertz. 2022. Evaluation of a semantic search approach based on AMR for information retrieval in image exploitation. In *2022 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, pages 1–6.

Tahira Naseem, Austin Blodgett, Sadhana Kumaravel, Tim O'Gorman, Young-Suk Lee, Jeffrey Flanigan, Ramón Astudillo, Radu Florian, Salim Roukos, and Nathan Schneider. 2022. DocAMR: Multi-sentence AMR representation and evaluation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3496–3505, Seattle, United States. Association for Computational Linguistics.

Marcel Nawrath, Agnieszka Nowak, Tristan Ratz, Danilo Walenta, Juri Opitz, Leonardo Ribeiro, João Sedoc, Daniel Deutsch, Simon Mille, Yixin Liu, Sebastian Gehrmann, Lining Zhang, Saad Mahamood, Miruna Clinciu, Khyathi Chandu, and Yufang Hou. 2024. On the role of summary content units in text summarization evaluation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 272–281, Mexico City, Mexico. Association for Computational Linguistics.

Ani Nenkova and Rebecca Passonneau. 2004. Evaluating content selection in summarization: The pyramid method. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 145–152, Boston, Massachusetts, USA. Association for Computational Linguistics.

Antonio MS Almeida Neto, Helena M Caseli, and Tiago A Almeida. 2020. Dense captioning using abstract meaning representation. In *Intelligent Systems: 9th Brazilian Conference, BRACIS 2020, Rio Grande, Brazil, October 20–23, 2020, Proceedings, Part I 9*, pages 450–465. Springer.

Long HB Nguyen, Viet H Pham, and Dien Dinh. 2021. Improving neural machine translation with AMR semantic graphs. *Mathematical Problems in Engineering*, 2021.

Takuro Ogawa and Ryosuke Saga. 2023. Inductive model using abstract meaning representation for text classification via graph neural networks. In *International Conference on Human-Computer Interaction*, pages 258–271. Springer.

Tim O'Gorman, Michael Regan, Kira Griffitt, Ulf Hermjakob, Kevin Knight, and Martha Palmer. 2018. AMR beyond the sentence: the multi-sentence AMR corpus. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3693–3702, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Dongsuk Oh, Jungwoo Lim, Kinam Park, and Heuiseok Lim. 2022. Semantic representation using subsymbolic knowledge in commonsense reasoning. *Applied Sciences*, 12(18).

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.

Juri Opitz, Angel Daza, and Anette Frank. 2021a. Weisfeiler-leman in the bamboo: Novel AMR graph metrics and a benchmark for AMR graph similarity. *Transactions of the Association for Computational Linguistics*, 9:1425–1441.

Juri Opitz and Anette Frank. 2021. Towards a decomposable metric for explainable evaluation of text generation from AMR. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1504–1518, Online. Association for Computational Linguistics.

Juri Opitz and Anette Frank. 2022a. Better Smatch = better parser? AMR evaluation is not so simple anymore. In *Proceedings of the 3rd Workshop on Evaluation and Comparison of NLP Systems*, pages 32–43, Online. Association for Computational Linguistics.

Juri Opitz and Anette Frank. 2022b. SBERT studies meaning representations: Decomposing sentence embeddings into explainable semantic features. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 625–638, Online only. Association for Computational Linguistics.

Juri Opitz, Philipp Heinisch, Philipp Wiesenbach, Philipp Cimiano, and Anette Frank. 2021b. Explainable unsupervised argument similarity rating with Abstract Meaning Representation and conclusion generation. In *Proceedings of the 8th Workshop on Argument Mining*, pages 24–35, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Juri Opitz, Letitia Parcalabescu, and Anette Frank. 2020. AMR similarity metrics from principles. *Transactions of the Association for Computational Linguistics*, 8:522–538.

Juri Opitz, Shira Wein, Julius Steen, Anette Frank, and Nathan Schneider. 2023. AMR4NLI: Interpretable and robust NLI measures from semantic graphs. In *Proceedings of the 15th International Conference on Computational Semantics*, pages 275–283, Nancy, France. Association for Computational Linguistics.

Elif Oral, Ali Acar, and Gülşen Eryiğit. 2022. Abstract meaning representation of Turkish. *Natural Language Engineering*, pages 1–30.

Christoph Otto, Jonas Groschwitz, Alexander Koller, Xiulin Yang, and Lucia Donatelli. 2024. A corpus of German Abstract Meaning Representation (DeAMR). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 286–292, Torino, Italia. ELRA and ICCL.

Martha Palmer, Paul Kingsbury, and Daniel Gildea. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31:71–106.

Xiaoman Pan, Taylor Cassidy, Ulf Hermjakob, Heng Ji, and Kevin Knight. 2015. Unsupervised entity linking with Abstract Meaning Representation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1130–1139, Denver, Colorado. Association for Computational Linguistics.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Jinwoo Park, Hosoo Shin, Dahee Jeong, and Junyeong Kim. 2024. Improving the representation of sentences with reinforcement learning and AMR graph. In *2024 IEEE International Conference on Consumer Electronics (ICCE)*, pages 1–4. IEEE.

Siyana Pavlova, Maxime Amblard, and Bruno Guillaume. 2023. Structural and global features for comparing semantic representation formalisms. In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 1–12, Nancy, France. Association for Computational Linguistics.

Hai Pham, Isma Hadji, Xinnuo Xu, Ziedune Degutyte, Jay Rainey, Evangelos Kazakos, Afsaneh Fazly, Georgios Tzimiropoulos, and Brais Martinez. 2024. Graph guided question answer generation for procedural question-answering. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2501–2525, St. Julian's, Malta. 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.

Jakob Prange, Nathan Schneider, and Lingpeng Kong. 2022. Linguistic frameworks go toe-to-toe at neuro-symbolic language modeling. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4375–4391, Seattle, United States. Association for Computational Linguistics.

Haoyi Qiu, Kung-Hsiang Huang, Jingnong Qu, and Nanyun Peng. 2024. AMRFact: Enhancing summarization factuality evaluation with AMR-driven negative samples generation. 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 594–608, Mexico City, Mexico. Association for Computational Linguistics.

Geetanjali Rakshit and Jeffrey Flanigan. 2021. ASQ: Automatically generating question-answer pairs using AMRs.

Sudha Rao, Daniel Marcu, Kevin Knight, and Hal Daumé III. 2017. Biomedical event extraction using Abstract Meaning Representation. In *BioNLP 2017*, pages 126–135, Vancouver, Canada,. Association for Computational Linguistics.

Michael Regan, Shira Wein, George Baker, and Emilio Monti. 2024. MASSIVE multilingual Abstract Meaning Representation: A dataset and baselines for hallucination detection. In *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics (*SEM 2024)*, pages 1–17, Mexico City, Mexico. Association for Computational Linguistics.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.

Leonardo F. R. Ribeiro, Mengwen Liu, Iryna Gurevych, Markus Dreyer, and Mohit Bansal. 2022. FactGraph: Evaluating factuality in summarization with semantic graph representations. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3238–3253, Seattle, United States. Association for Computational Linguistics.

Leonardo F. R. Ribeiro, Jonas Pfeiffer, Yue Zhang, and Iryna Gurevych. 2021a. Smelting gold and silver for improved multilingual AMR-to-Text generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 742–750, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Leonardo F. R. Ribeiro, Martin Schmitt, Hinrich Schütze, and Iryna Gurevych. 2021b. Investigating pretrained language models for graph-to-text generation. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 211–227, Online. Association for Computational Linguistics.

Mrinmaya Sachan and Eric Xing. 2016. Machine comprehension using rich semantic representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 486–492, Berlin, Germany. Association for Computational Linguistics.

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.

Yuichiro Sawai, Hiroyuki Shindo, and Yuji Matsumoto. 2015. Semantic structure analysis of noun phrases using Abstract Meaning Representation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 851–856, Beijing, China. Association for Computational Linguistics.

Nikolaus Schrack, Ruixiang Cui, Hugo López, and Daniel Hershcovich. 2022. Can AMR assist legal and logical reasoning? In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1555–1568, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Verena Severina and Masayu Leylia Khodra. 2019. Multidocument abstractive summarization using abstract meaning representation for indonesian language. In *2019 International Conference of Advanced Informatics: Concepts, Theory and Applications (ICAICTA)*, pages 1–6.

Kaize Shi, Xueyao Sun, Li He, Dingxian Wang, Qing Li, and Guandong Xu. 2023. AMR-TST: Abstract Meaning Representation-based text style transfer. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4231–4243, Toronto, Canada. Association for Computational Linguistics.

Kaize Shi, Xueyao Sun, Qing Li, and Guandong Xu. 2024. Compressing long context for enhancing RAG with AMR-based concept distillation.

Mollie Shichman, Claire Bonial, Austin Blodgett, Taylor Hudson, Francis Ferraro, and Rachel Rudinger. 2023. Use defines possibilities: Reasoning about object function to interpret and execute robot instructions. In *Proceedings of the 15th International Conference on Computational Semantics*, pages 284–292, Nancy, France. Association for Computational Linguistics.

Kanchan Shivashankar, Khaoula Benmaarouf, and Nadine Steinmetz. 2022. From graph to graph: Amr to sparql.

Ziyi Shou, Yuxin Jiang, and Fangzhen Lin. 2022. AMR-DA: Data augmentation by Abstract Meaning Representation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3082–3098, Dublin, Ireland. Association for Computational Linguistics.

Ziyi Shou and Fangzhen Lin. 2023. Evaluate AMR graph similarity via self-supervised learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16112–16123, Toronto, Canada. Association for Computational Linguistics.

Marco Antonio Sobrevilla Cabezudo and Thiago Pardo. 2019. Towards a general Abstract Meaning Representation corpus for Brazilian Portuguese. In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 236–244, Florence, Italy. Association for Computational Linguistics.

Linfeng Song, Daniel Gildea, Yue Zhang, Zhiguo Wang, and Jinsong Su. 2019. Semantic neural machine translation using AMR. *Transactions of the Association for Computational Linguistics*, 7:19–31.

Linfeng Song, Ante Wang, Jinsong Su, Yue Zhang, Kun Xu, Yubin Ge, and Dong Yu. 2020. Structural information preserving for graph-to-text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7987–7998, Online. Association for Computational Linguistics.

Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. A graph-to-sequence model for AMR-to-text generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1616–1626, Melbourne, Australia. Association for Computational Linguistics.

Katharina Stein, Lucia Donatelli, and Alexander Koller. 2023. From sentence to action: Splitting AMR graphs for recipe instructions. In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 52–67, Nancy, France. Association for Computational Linguistics.

Nadine Steinmetz. 2023. Entity linking for KGQA using AMR graphs. In *European Semantic Web Conference*, pages 122–138. Springer.

Reza Takhshid, Razieh Shojaei, Zahra Azin, and Mohammad Bahrani. 2022. Persian Abstract Meaning Representation. *arXiv preprint arXiv:2205.07712*.

Christopher Tam, Richard Brutti, Kenneth Lai, and James Pustejovsky. 2023. Annotating situated actions in dialogue. In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 45–51, Nancy, France. Association for Computational Linguistics.

Jingxuan Tu, Timothy Obiso, Bingyang Ye, Kyeongmin Rim, Keer Xu, Liulu Yue, Susan Windisch Brown, Martha Palmer, and James Pustejovsky. 2024. GLAMR: Augmenting AMR with GL-VerbNet event structure. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7746–7759, Torino, Italy. ELRA and ICCL.

Zdeňka Urešová, Jan Hajič, and Ondřej Bojar. 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.

Jens E. Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O'Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, Jan Hajič, James H. Martin, Stephan Oepen, Martha Palmer, James Pustejovsky, Rosa Vallejos, and Nianwen Xue. 2021. Designing a uniform meaning representation for natural language processing. *KI - Künstliche Intelligenz*.

Pavlo Vasylenko, Pere Lluís Huguet 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.

Martin Verrev. 2023. Evaluation of semantic parsing frameworks for automated knowledge base construction. In *Intelligent Systems Design and Applications*, pages 554–563, Cham. Springer Nature Switzerland.

Trong Sinh Vu and Minh Le Nguyen. 2019. An empirical evaluation of AMR parsing for legal documents. In *New Frontiers in Artificial Intelligence: JSAI-isAI 2018 Workshops, JURISIN, AI-Biz, SKL, LENLS, IDAA, Yokohama, Japan, November 12–14, 2018, Revised Selected Papers*, pages 131–145. Springer.

Trong Sinh Vu, Minh Le Nguyen, and Ken Satoh. 2022. Abstract meaning representation for legal documents: an empirical research on a human-annotated dataset. *Artificial Intelligence and Law*, 30(2):221–243.

Cunxiang Wang, Zhikun Xu, Qipeng Guo, Xiangkun Hu, Xuefeng Bai, Zheng Zhang, and Yue Zhang. 2023. Exploiting Abstract Meaning Representation for open-domain question answering. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2083–2096, Toronto, Canada. Association for Computational Linguistics.

Yanshan Wang, Sijia Liu, Majid Rastegar-Mojarad, Liwei Wang, Feichen Shen, Fei Liu, and Hongfang Liu. 2017. Dependency and AMR embeddings for drug-drug interaction extraction from biomedical literature. In *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology,and Health Informatics*, ACM-BCB '17, page 36–43, New York, NY, USA. Association for Computing Machinery.

Shira Wein, Lucia Donatelli, Ethan Ricker, Calvin Engstrom, Alex Nelson, Leonie Harter, and Nathan Schneider. 2022. Spanish Abstract Meaning Representation: Annotation of a general corpus. In *Northern European Journal of Language Technology, Volume 8*, Copenhagen, Denmark. Northern European Association of Language Technology.

Shira Wein and Nathan Schneider. 2022. Accounting for language effect in the evaluation of cross-lingual AMR parsers. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3824–3834, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Shira Wein and Nathan Schneider. 2024a. Assessing the Cross-linguistic Utility of Abstract Meaning Representation. *Computational Linguistics*, pages 1–55.

Shira Wein and Nathan Schneider. 2024b. Lost in translationese? reducing translation effect using Abstract Meaning Representation. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 753–765, St. Julian's, Malta. Association for Computational Linguistics.

Shira Wein, Zhuxin Wang, and Nathan Schneider. 2023. Measuring fine-grained semantic equivalence with Abstract Meaning Representation. In *Proceedings of the 15th International Conference on Computational Semantics*, pages 144–154, Nancy, France. Association for Computational Linguistics.

Runxin Xu, Peiyi Wang, Tianyu Liu, Shuang Zeng, Baobao Chang, and Zhifang Sui. 2022. A two-stream AMR-enhanced model for document-level event argument extraction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5025–5036, Seattle, United States. 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.

Zhiyang Xu, Jay Yoon Lee, and Lifu Huang. 2023. Learning from a friend: Improving event extraction via self-training with feedback from Abstract Meaning Representation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10421–10437, Toronto, Canada. 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).

Bohao Yang, Kun Zhao, Chen Tang, Liang Zhan, and Chenghua Lin. 2024. Structured information matters: Incorporating abstract meaning representation into llms for improved open-domain dialogue evaluation. *arXiv preprint arXiv:2404.01129*.

Yuqing Yang, Qipeng Guo, Xiangkun Hu, Yue Zhang, Xipeng Qiu, and Zheng Zhang. 2023. An AMR-based link prediction approach for document-level event argument extraction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12876–12889, Toronto, Canada. Association for Computational Linguistics.

Shaowei Yao, Tianming Wang, and Xiaojun Wan. 2020. Heterogeneous graph transformer for graph-to-sequence learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7145–7154, Online. Association for Computational Linguistics.

Dongran Yu, Bo Yang, Dayou Liu, Hui Wang, and Shirui Pan. 2023. A survey on neural-symbolic learning systems. *Neural Networks*, 166:105–126.

Zixuan Zhang and Heng Ji. 2021. Abstract Meaning Representation guided graph encoding and decoding for joint information extraction. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 39–49, Online. Association for Computational Linguistics.

Zixuan Zhang, Nikolaus Parulian, Heng Ji, Ahmed Elsayed, Skatje Myers, and Martha Palmer. 2021. Fine-grained information extraction from biomedical literature based on knowledge-enriched Abstract Meaning Representation. 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 6261–6270, Online. Association for Computational Linguistics.

Qihui Zhao, Tianhan Gao, and Nan Guo. 2023. Document-level relation extraction based on sememe knowledge-enhanced abstract meaning representation and reasoning. *Complex & Intelligent Systems*, 9(6):6553–6566.

A Reification and Levi Graph

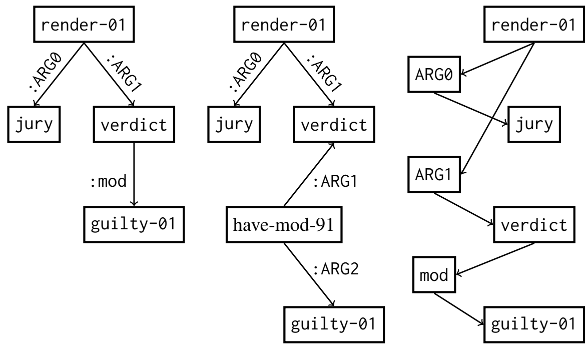


```
(r / render-01
    :ARG0 (j / jury)
    :ARG1 (v / verdict
        :mod (g / guilty-01)))

(r / render-01
    :ARG0 (j / jury)
    :ARG1 (v / verdict
        :ARG1-of (h / have-mod-91
             :ARG2 (g / guilty-01))))

(r / render-01
    (a1 / ARG0
         (j / jury))
    (a2 / ARG1
         (v / verdict
              (m / mod
                   (g / guilty)))))
```

Figure 3: Three equivalency-preserving AMR transformations for "The jury rendered a guilty verdict." Left/-Top: Standard AMR. Middle/Middle: Reification with AMR rules. Right/Bottom: Bipartite Levi Graph with unlabeled edges. While Levi Graphs are not AMR-specific, reification is. Per the AMR guidelines (Banarescu et al., 2019), any labeled edge not in a standardized set (`:opN`, `:argN`, etc.) is generalized to a new structure, where the old edge assumes the position of a node linked with `:opN`/`:argN` to the original structure (in the example, `:mod` triggers the node `have-mod-91` with arguments `:ARG1` and `:ARG2`).