A Survey of Meaning Representation – From Theory to Practical Utility

Anonymous ACL submission

Abstract

Symbolic meaning representations of natural language text have been studied since at least the 1960s. With the availability of large annotated corpora, and more powerful machine learning tools, the field has recently seen several new developments. In this survey, we study today's most prominent Meaning Representations Frameworks. We shed light on both their theoretical properties and on their practical research environment, which includes datasets, parsers, applications, and future challenges.

1 Introduction

Being able to represent the semantic structure of a text has been an important research goal since the early days of NLP. Early works started to develop natural language interfaces on specific databases by transforming raw text into an executable language, using formalisms such as SQL, first-order logic or lambda-calculus (Mooney, 1996; Wong and Mooney, 2006; Mooney, 2007). Another avenue of research, which is the focus of this work, has developed general-purpose, non-executable Meaning Representations (MRs), inspired by formal grammars. These often take the form of human-readable graphs. Figure 1 shows an example.

Such MRs are used to improve the accuracy of NLP systems in tasks such as summarization or machine translation (Gao and Vogel, 2011; Liu et al., 2015; Mohamed and Oussalah, 2019; Liao et al., 2018; Song et al., 2019; Ribeiro et al., 2022). In the age of large language models (LLMs), they also get leveraged for their interpretability, e.g., to enhance semantic search (Bonial et al., 2020; Cai et al., 2022; Opitz and Frank, 2022b) or natural language inference (Opitz et al., 2023b). They are also used to generate paraphrases (Cai et al., 2021), augment training data (Shou et al., 2022), or to do style-transfer (Jangra et al., 2022; Shi et al., 2023).

In this survey, we provide a structured overview of current Meaning Representations Formalisms. Several other surveys have discussed MRs before us. However, they are either focused on linguistic theory (Abend and Rappoport, 2017; Pavlova et al., 2023) and thus tend to neglect applications, parsers, and resources, or they focus on the practical application (Verrev, 2023), and present the different formalisms only in a few lines. Our survey takes a balanced stance: It presents both the different formalisms and their applications, resources, and parsers. This balance allows us to describe a bigger picture and outline commonalities and open challenges. Our survey thus aims to be a handy reference for anyone who wishes to choose, understand, build, or use a Meaning Representation.

In Section 2 we introduce the main concepts and properties of MRs. Section 3 tackles Shallow MRs, and Section 4 Deep MRs. Finally, Section 5 discusses open challenges in the domain.

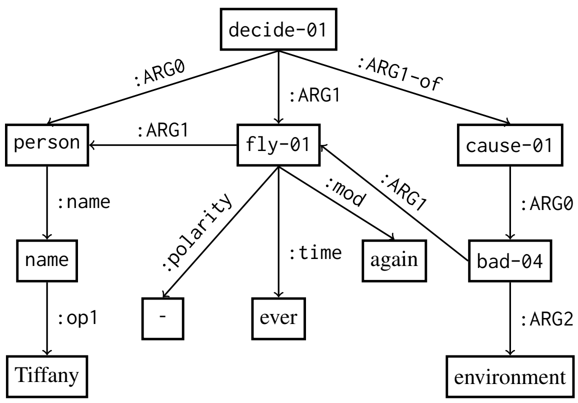

Figure 1: AMR graph for the sentence "Tiffany decided that she would never fly again, because it is bad for the environment".

2 Meaning Representations

Given a text in natural language, MR parsing is the task of producing a symbolic representation of its meaning, as it is understood by a language speaker (Abend and Rappoport, 2017). Different Meaning Representation Formalisms (MRFs) have

MRF	Subevents	Shape	Compositionality	Node type (Flavor)	Edge type
SRL	✗	Tree	Non-Compositional	Text spans (1)	Numbered
RST	✗	Tree	Compositional	Text spans (1)	Theory-oriented
UDS	✗	Tree	Compositional	Text spans (1)	Numbered & Theory-oriented
SDP	✓	Graph	Non-Compositional	Augmented Word Spans (0)	Numbered
EDS	✓	Graph	Non-Compositional	Augmented Text Spans (1)	Numbered
UCCA	✓	Tree	Compositional	Text spans (1)	Theory-oriented
AMR	✓	Graph	Non-Compositional	Synsets (Propbank) (2)	Predicate-dependent
DRS	✓	Graph	Compositional	Synsets (WordNet) (2)	Predicate-independent

Table 1: Properties of the Meaning Representation Frameworks that we survey. The middle line separates shallow and deep formalisms.

been developed. Here, we focus on graph-like MRFs that target English language sentences. Figure 1 shows a MR in an MRF called AMR.

MRFs often adopt Neo-Davidsonian semantics, and see *events* as the central elements of sentences. These events center around a *predicate*, which indicates the type of the event, and is most often a verb (`decide-01` or `fly-01` in Figure 1). The *arguments* correspond to the entities that participate in the event ("Tiffany" in the example), or to the circumstances of the event, such as place or manner (a negative polarity "-", in our example). The *semantic role* specifies the participant's role in the event, e.g., the semantic role of Tiffany in the `decide-01` event is the subject (the `ARG0` in AMR jargon).

The information in a MR can be decomposed into three levels, with the event level in the middle. On the sub-event level, the arguments of an event can themselves be decomposed into more atomic components. In our example, "bad for the environment" is modeled by the link from `bad-4` to `environment` with the semantic role `ARG2`. On the supra-event level, events can also be linked, using *discourse relations*. In our example, the `cause-01` node connects `decide-01` and `bad-04`, meaning that the decision was taken *because* flying is bad for the environment. Discourse relations can also link events across sentence borders.

Different MRFs vary these general ideas along several axes, which we show in Table 1. First, not all MRFs can represent sub-events (Column 2 in Table 1), so we call a MRF *deep* if it represents sub-events, and *shallow* otherwise. Second, MRFs construct either trees (where each node has at most one parent) or full-fledged graphs (Column 3). Figure 1 is a full-fledged graph: `fly-01` and `person` play two different roles, and participate in a cycle. Third, some MRFs are *compositional* (Column 4), which means that they create nodes that compose the meaning of other nodes. Our example in Figure 1 is not compositional: every node corresponds to one element. However, we can imagine creating a node that represents the fact that the `fly-01` event is negated. This would then be a compositional node.

MRFs can further be distinguished by their types of nodes (Column 5): Nodes can be labeled with a span from the text, but they can also be augmented with extra information such as a POS tag or other information. Some representations even use abstractions such as synsets from predefined vocabularies, to help reduce (or even eliminate) lexical ambiguity, and make events invariant to surface form. The node type is closely related to the Flavor hierarchy proposed by Oepen et al. (2019). It differentiates Meaning Representations based on anchoring, i.e. on the explicit correspondence between nodes and the input sentence. Flavor 0 means that each node injectively corresponds to one word, while Flavor 1 relaxes the anchoring constraints, allowing a node to correspond to a whole span, and the same span to correspond to several nodes, and Flavor 2 marks the absence of explicit node-text links.

Finally, the MRFs differ in their edge type (Column 6): Some MRFs use roles that depend on a specific linguistic theory, like `elaboration` (discourse theory) or `scene` (cognitive science). These schemes can describe only a limited array of relations, and for instance do not distinguish agents and patients. Other representations are more specific and use numbered semantic roles (`A0`, `A1`, ...). In these schemes, `A0` and `A1` usually correspond to Dowty's Proto-Agent and Proto-Patient (Dowty, 1991), respectively. These proto-roles are defined by their features: Typical agent features are *awareness, movement*, and *volition*, while typical patient features are *change of state*, *being stationary*, etc. The other semantic roles (`A2`, `A3`, ...) usually do not have such a predefined meaning. Again, other MRFs are more specific, and use predicate-independent semantic roles that distinguish finer roles such as `Agent` and `Patient`. Finally, some MRFs make the meaning of the role dependent on the predicate: in Figure 1 `ARG0` means "pilot, agentive entity capable of flight" for `fly-01`, while it means "decider" for `decide-01`. These MRFs thus describe their arguments very specifically.

3 Shallow Meaning Representation Frameworks

3.1 Semantic Roles

A prominent Shallow Meaning Representation Framework is Semantic Roles (SR, Gildea and Jurafsky, 2000), also known as Semantic Role Labeling (SRL). Given an input sentence and a predicate, its purpose is to determine the predicate's arguments and their semantic roles. It focuses on event-level relations, which means that predicates are verbs. There are (at least) three different implementations of semantic roles. The most popular one is *PropBank SRL*, where semantic roles are split into core and non-core roles according to PropBank (Palmer et al., 2005). The non-core roles are also called modifiers, and they always have the same meaning: `ARGM-CAU` indicates cause, `ARGM-LOC` indicates location, etc. The meaning of core roles ($\texttt{ARG}_{2...n}$) depends on the predicate. However, `ARG0` and `ARG1` usually correspond to Dowty's Proto-Agent and Proto-Patient (Dowty, 1991). Other paradigms exist: *FrameNet* (Baker et al., 1998) *SRL* generalizes descriptions across similar verbs (e.g., say, speak) as well as similar nouns and other words (e.g., speech). *Semantic Proto-Role Labeling* (SPRL) aims at directly approximating Dowty's agent and patient roles with features such as movement, awareness, etc.

Figure 2 shows a merger of three parsings for our example (in PropBank-SRL style), for the predicates "decided", "fly", and "is". Having only one predicate node and its arguments, an SRL graph is a dependency tree. No node abstraction is performed, meaning all nodes are text spans.

SRL is a rather light annotation, and it is used to enhance LLMs (Zhang et al., 2020b) for downstream tasks such as Fact Checking (Zhong et al., 2020), Question Answering (Pillai et al., 2018), and Summarization (Mohamed and Oussalah, 2019).

Resources. (PropBank-)SRL has been the focus of several shared tasks, which provided datasets that are used to this day. CoNLL 2005 (Carreras and Màrquez, 2004, 2005) introduced span-based SRL, while CoNLL 2008 (Surdeanu et al., 2008) and 2009 (Hajič et al., 2009) introduced dependency-based SRL (which labels only syntactic heads of arguments). Other datasets propvide FrameNet SRL (Burchardt and Pennacchiotti, 2008; Das and Smith, 2011; Hartmann et al., 2017) and SPRL annotations (Reisinger et al., 2015; White et al., 2016a).

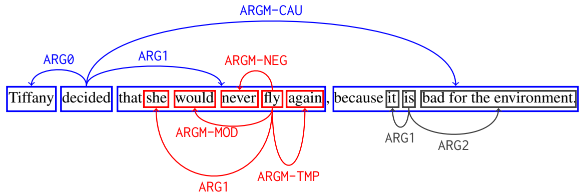

Figure 2: Semantic Role Labeling of our example sentence in span-graph style.

Parsing. Regardless of the flavor of SRL, many approaches for parsing (or labeling) are heavily reliant on syntactic features (Pradhan et al., 2005; Punyakanok et al., 2008; Li et al., 2018; Fei et al., 2021). The progress in Neural Networks allowed systems to become more syntax-agnostic (Zhou and Xu, 2015; He et al., 2017; Tan et al., 2018; Rudinger et al., 2018; Arora et al., 2022; Spaulding et al., 2023), so much that recent approaches extract not just the arguments, but also the predicates themselves (Cai et al., 2018; He et al., 2018; Zhang et al., 2021), which is particularly appealing in the scope of Meaning Representation.

3.2 Rhetorical Structure Theory

Rhetorical Structure Theory (RST, Mann and Thompson, 1988) takes interest in discourse relations. It sees the text as a sequence of *Elementary Discourse Units (EDUs)*, which roughly correspond to events, and seeks to identify the relations between these units, such as `Condition`, `Contrast`, `Cause`, `Result`, or `Elaboration`. RST models a text as a tree, in which discourse relations are recursively applied connect discourse units. Leaf nodes are EDUs (text spans), while inner nodes are unlabeled, and represent sets of EDUs. Figure 3 shows the RST MR of our example sentence. We see that EDUs correspond to events, as they coincide with the spans delimited by predicates and arguments in the SRL graph. Each discourse relation links a *satellite* (supporting information) to a *nucleus* (central information). In our example, the nucleus of the `Reason` relation is the fact that Tiffany decided to never fly again, and the satellite is the reason for that decision. The repertory of discourse relations depends on the dataset.

Discourse relations can cross sentence boundaries, which means that one rhetorical structure can represent a multi-sentence document. RST is used for summarization (Xu et al., 2020) and question answering (Ouyang et al., 2021), and is even used for argument mining (Peldszus and Stede, 2013; Mitrović et al., 2017; Chakrabarty et al., 2019).

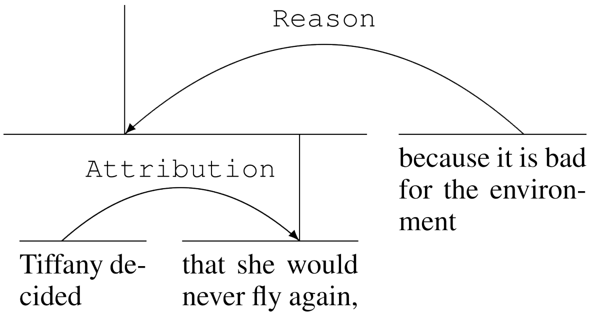

Figure 3: RST-DT style annotation for our example.

Resources. The main dataset for RST is RST-DT (Carlson et al., 2001), which defines 78 discourse relations, divided into 16 classes. The dataset contains 385 documents from the Wall Street Journal corpus, with around 20.000 EDUs.

Parsing is usually performed in two steps: EDU Segmentation and Tree Building. Wang et al. (2018) achieves a 95% F1-score on segmentation with a Bi-LSTM-CRF-based model, while human performance is only marginally better, at 98%. First approaches for Tree Building (Soricut and Marcu, 2003; Hernault et al., 2010) used hand-crafted features. Ji and Eisenstein (2014) introduced the first RST-DT neural parser, followed by bottom-up parsers (Li et al., 2016; Braud et al., 2017; Wang et al., 2017; Yu et al., 2018), and more recently top-down ones (Lin et al., 2019; Zhang et al., 2020a; Kobayashi et al., 2020). Though they have different approaches, Nguyen et al. (2021) and Koto et al. (2021) define the state-of-the-art with Parseval scores (Morey et al., 2017) of 50.2 and 50.3, respectively (human performance is at 55.0).

3.3 Universal Decompositional Semantics

Universal Decompositional Semantics (UDS) is a multi-layer semantic annotation scheme, which means that it allows annotating the same sentence on different dimensions. These dimensions are, e.g., factuality and time for predicates, or genericity and word sense for arguments. UDS builds a semantic compositional tree, where leaf nodes are words of the sentence (or special tokens) and inner nodes represent larger semantic units. The graph structure is based on PredPatt (White et al., 2016a), a pattern-based framework for predicate-argument extraction that operates on (syntactic) Universal Dependencies (UD, de Marneffe et al., 2021). It focuses on event-level relations, which means the extracted structure is close to that of merged SRL graphs. UDS uses Dowty's Proto Roles, which, as described above, describe features of event participants and how they are affected by the event (*movement*, *volition*, *change of state*, and so on).

Resources. The UDS dataset can be accessed through the Decomp Toolkit (White et al., 2020). The original annotations include proto-roles (Reisinger et al., 2015), word sense, and factuality. They were augmented with annotations on time (Vashishtha et al., 2019), and generalizing statements (Govindarajan et al., 2019), and also some discourse relations (Gantt et al., 2022).

Parsing. UDS Parsing is a fairly unexplored task. Zhang et al. (2018) performs cross-lingual UDS parsing with a pipeline approach performing graph transduction, coreference resolution and semantic proto-role labeling. Stengel-Eskin et al. (2020) proposes an end-to-end parser with an encoder-decoder structure, while Stengel-Eskin et al. (2021) parses UD and UDS jointly.

4 Deep Meaning Representation Frameworks

Deep Meaning Representation Frameworks go further than shallow ones by representing relations at all levels of the text, in particular at the sub-event level. They aim to model the meaning of the text exhaustively, representing as many phenomena as possible (noun phrases, negations, comparisons, modifiers, time, cause, etc.).

4.1 Semantic Dependencies

Semantic Dependencies (SD), also known as Semantic Dependency Parsing (SDP), is a family of MR frameworks that are based on the SemEval 2014 & 2015 challenges (Oepen et al., 2014, 2015). Their aim is to go further than syntactic dependency parsing, and to represent the semantic structure of a sentence. Four main frameworks have been proposed, derived from independent annotation schemes with different formalisms: DM

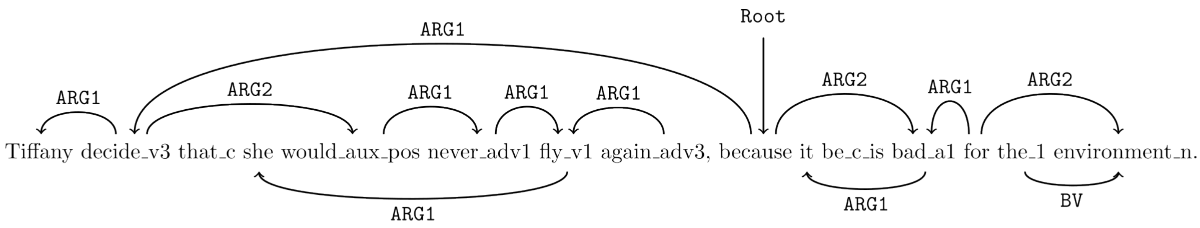

Figure 4: A Semantic Dependency Parse in DM-style for our running example.

(DELPH-IN MRS-Derived Bi-Lexical Dependencies, Flickinger et al., 2012), PAS (Enju Predicate-Argument Structures, Miyao, 2006), PSD (Prague Semantic Dependencies, Hajič et al., 2012), and CCD (Combinatory Categorial Grammar Dependencies, Hockenmaier and Steedman, 2007).

All frameworks see the semantic structure as a bi-lexical dependency graph: every node corresponds to exactly one word in the sentence. In contrast to syntactic dependency trees, a graph structure is required, as nodes can have several incoming edges (a phenomenon called *re-entrancy*) if a word is the argument of several predicates, as well as none if they are semantically vacuous. A node is a word that can be augmented with its lemma, POS-tag or framework-specific frame. The exact vocabulary of semantic roles, as well as the way the graph models different phenomena, varies across frameworks. Most of them use unspecific semantic roles (`ARG1`, `ARG2`, `ARG3`, ...). Nevertheless, similar to SRL, `ARG1` and `ARG2` usually correspond to Dowty's Proto-Agent and Proto-Patient.

Still, SDP has the advantage to be easily understandable by human readers. Figure 4 shows DM annotations for our example sentence. We see that relations go all the way to the token level: the noun phrase "bad for the environment" is seen as an object of interest, with "for" being a predicate, with the arguments "bad" and "environment".

Resources. Oepen et al. (2016) proposes a corpus with annotations for all four frameworks, with close to 37.000 English sentences. The dataset also provides a corpus of PAS annotations on Chinese text, and PSD annotations on Czech text.

Parsing. Most parsing approaches for SDP are inspired by syntactic dependency parsing (Dozat and Manning, 2018; Fernández-González and Gómez-Rodríguez, 2020). ACE (Wang et al., 2021) achieves state of the art results in SDP (on DM, PSD and PAS) and other structured prediction tasks by acting on embeddings concatenation.

Variations. English Resource Grammar (ERG), which DM is a reduction of, produces MRs in the Minimal Recursion Semantics (Copestake et al., 2005). These structures are particularly expressive and can model scope, but they are also complex to read and exploit. Elementary Discourse Structures (EDS, Oepen and Lønning, 2006) try to reduce this complexity by making the graph non-compositional. The main difference between EDS and DM is that EDS are Flavor 1 graphs with stronger node abstraction: in addition to POS tags and identifiers, nodes can be labeled with properties, such as time or number.

4.2 Universal Conceptual Cognitive Annotation

The Universal Conceptual Cognitive Annotation (UCCA, Abend and Rappoport, 2013) is a semantic annotation scheme aiming to be "universal", hence it wants to be resistant to syntactic variation within and across languages. An UCCA Representation takes the form of a compositional tree whose leaf nodes are the words of the sentence, and intermediate nodes, called `units`, are unlabeled. UCCA identifies 3 levels of semantic information. On the central level, scene units correspond to events. They are linked to a predicate, to its core arguments by a generic label `participant`, as well as to non-core arguments using several other labels (see Figure 5). On the lower level, *sub-scene units* help specify the participants of a scene. Finally, *superparallel units* can link two scenes with generic `parallel scene` edges, and possibly a cue word indicating the type of discourse relation with a `linker` edge. At any level, *functional units* can represent phenomena such as prepositions, articles, or expletive pronouns. UCCA can annotate several sentences in a single graph.

There are very few semantic roles in UCCA, which makes the annotation task more accessible to non-experts and portable to other languages. Semantic roles have a generic interpretability, but it can be hard to exploit them directly: for instance, the `participant` role doesn't make a difference between what would be labeled as `ARG0` (Agent) and `ARG1` (Patient) in other frameworks. UCCA is multi-layered, which makes it possible to add extensions to the representation, for instance to annotate co-reference links, more specific semantic roles, or more abstract node types. UCCA is cross-lingual, and as such found applications in machine translation (Slobodkin et al., 2022; Birch et al., 2016), but also in text simplification (Sulem et al., 2018a,b).

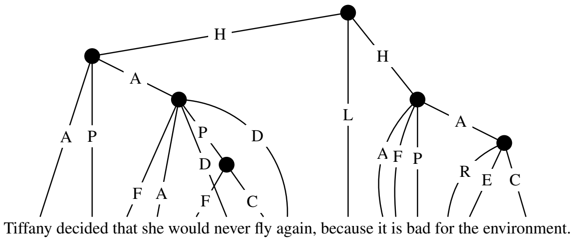

Figure 5: UCCA graph for our example. H: Parallel Scene, L: Linker, P: Process, A: Participant, D: Adverbial, F: Function C: Center, E: Elaborator, R: Relator.

Resources. UCCA comes with a large annotated multilingual corpus (Abend and Rappoport, 2013). Its English version includes annotations of Wikipedia, the Web- and Penn Treebank, *Twenty Thousand Leagues Under the Sea*, and *The Little Prince*, with a total of 1350 passages (more than 200k tokens). Some of the sources were also annotated in French, German, Hebrew, and Russian.

Parsing. The first proposed parser for UCCA (Hershcovich et al., 2017) was transition-based. Other methods exploit constituency parsers (Jiang et al., 2019; Bölücü and Can, 2021). Nowadays, the best parsers are sequence-to-sequence models (Ozaki et al., 2020; Samuel and Straka, 2020).

4.3 Abstract Meaning Representation

Abstract Meaning Representation (AMR, Banarescu et al., 2013) aims at further abstracting away from syntax, even mapping named entities to Wikipedia. AMR has no explicit alignments between nodes and the text. The representation itself takes the form of a rooted, acyclic, directed dependency graph, where each node is labeled with a *concept*, and represents an instance of this concept. The root of an AMR is used for modeling the focus, or main event, of a text. Figure 1 shows the AMR graph for our running example.

AMR performs strong node abstraction: node labels can be PropBank frames for explicitly modeling entities, unambiguous English words, or special frames (e.g. for dates, modality, negation, comparisons, or family relationships). Semantic roles are either PropBank roles, which have accessible predicate-specific interpretation, or manually-crafted ones (e.g. `:name`, `:location`, `:cause`, `:concession`, `:month`, `:poss`, `degree`...).

Many AMR relations can be reified and used as concepts, thus allowing the focus to be on the relation itself. AMR also makes use of re-entrancy: in our example, Tiffany appears only once as a node, and is linked to both `decide-01` and `fly-01`. AMR also represents explicit quantities and temporal relations. This makes AMR graphs nearly unambiguous. However, the lack of explicit scope can still lead to ambiguity: in our example, it is unclear whether what is bad for the environment is only the node "fly", or the subgraph meaning "that Tiffany will never fly again" – which is the opposite of the actual meaning of the sentence.

Of all MRFs, AMR has probably garnered the most attention in recent years. It has been used in tasks such as machine translation (Song et al., 2019), question answering (Kapanipathi et al., 2021; Lim et al., 2020; Xu et al., 2021), toxic content detection (Elbasani and Kim, 2022), semantic search and natural language inference (Opitz and Frank, 2022b; Opitz et al., 2023b) and social reasoning (Chanin and Hunter, 2023).

Resources. The most important AMR corpus is the AMR Annotation Release (Banarescu et al., 2013). It was constructed fully manually, and contains 60.000 AMR graphs in its latest (3.0) version, including multi-sentence graphs (O'Gorman et al., 2018). AMR graphs are often linearized in the 'Penman' form (Kasper, 1989), which is easy to read, and allows processing with neural models in a sequence-to-sequence manner (the Penman uses a depth-first traversal and can, in principle, linearize any directed and rooted graph).

Parsing. Many AMR parsers have been proposed through the years, graph-based (Flanigan et al., 2014; Werling et al., 2015; Cai and Lam, 2020), transition-based (Wang et al., 2015; Vilares and Gómez-Rodríguez, 2018; Lee et al., 2020), or seq-2-seq (Barzdins and Gosko, 2016; Peng et al., 2018; Bevilacqua et al., 2021), possibly leveraging adapters to better incorporate graph topology (Vasylenko et al., 2023). Most systems of the 2020s leverage large pre-trained language models and achieve strong performance on AMR 3.0

Extensions. AMR has been extended to model tense and aspect (Donatelli et al., 2018), as well as scope (Pustejovsky et al., 2019), and larger documents (Naseem et al., 2021). The BabelNet Meaning Representation (Navigli et al., 2022) aims at making it multilingual by using BabelNet synsets for concepts (Navigli et al., 2021) and semantic roles from VerbAtlas (Di Fabio et al., 2019). Perhaps even more ambitious, the Universal Meaning Representation (UMR, Van Gysel et al., 2021) aims at compensating all main shortcomings of AMR, adding aspect and scope, integrating document-level annotations with coreference, temporal and modal relations between sentences, and making the representation language-agnostic.

4.4 Discourse Representation Structure

Discourse Representation Structure (DRS) is the fruit of Discourse Representation Theory (DRT, Kamp, 1981; Kamp and Reyle, 1993) and provides a meaning representation that fully integrates with first order logic. We focus here on the characteristics of the DRS format used in the Parallel Meaning Bank (PMB, Abzianidze et al., 2017), based on Segmented Discourse Representation Theory (Asher and Lascarides, 2003), which implements discourse relations. A Discourse Representation Structure (DRS) is not a graph, but a recursive structure of nested boxes. Figure 6 shows the representation of our example sentence.

As in AMR, elements are represented by a concept, here a Wordnet synset (Miller, 1995), accompanied by an identifier. Wordnet has a very wide coverage of English, which means that most labels are fully abstract. Semantic roles are taken from VerbNet (Kipper et al., 2000), augmented by handcrafted roles (e.g. `Quantity`, `Name`, `Owner`, `Time`), some of which are used specifically for comparison. Yet, these roles are generic, and no predicate-specific interpretation is available.

Usually, a simple box represents an event (similar to an EDU). Discourse relations are represented similarly to semantic roles, but with boxes as arguments. This means that DRS is compositional, and naturally equipped for multi-sentence representation. Modal logic operators can also be applied to boxes (negation, possibility, and necessity), which allows for a precise scoping of these operators: in the example, "she will never fly again" is represented as the negation of the box expressing that Tiffany flies at some point in the future.

Even though there is no ideal way to transform DRS into a graph (Abzianidze et al., 2020), we can see concepts as nodes, and semantic roles as labels of the edges between these nodes. Boxes would be another type of nodes, with discourse relations linking them. The most recent development of DRS, the Sequence Notation (Bos, 2023) proposes a similar graph equivalent. With this view, DRS are compositional graphs, with high-level nodes representing scope.

Resources. DRS annotations are hard to produce even for experts, which makes constructing large corpora difficult. The Groningen Meaning Bank (GMB, Basile et al., 2012) was the first DRS corpus, followed by the Parallel Meaning Bank (PMB, Abzianidze et al., 2017). These banks were built using an automatic pipeline using the rule-based parser Boxer (Bos, 2008). The PMB tries to make DRS language-neutral by associating English documents with translations to one or several languages. The latest release contains almost 10.000 "gold", i.e., human-checked, English documents.

Parsing. Several DRS parsers are available, exploiting transition-based parsing (Evang, 2019), DAG Grammars (Fancellu et al., 2019) or POS-tags and dependency graphs (van Noord, 2019). Modern parsers use LLMs (van Noord et al., 2018, 2020) and generally outperform older ones.

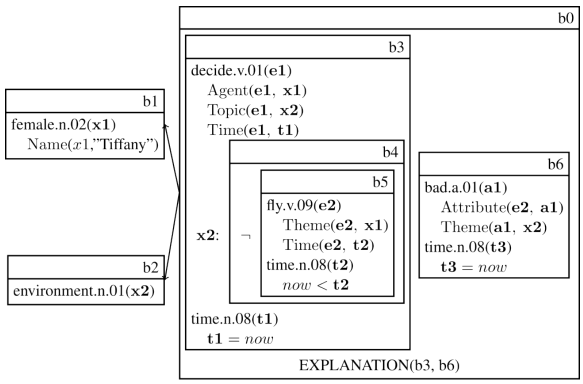

Figure 6: DRS for our running example

5 Current Research Trends

Synthesizing insights from our overview of MRFs, we see research challenges in three main areas: MR design, MR parsing, and MR applications.

5.1 Trends in MR design

MRs seem to lend themselves to multi-linguality, since they represent semantic concepts such as *agent*, *patient*, *instrument*, and *cause* that appear universal. However, MRs often have a strong flavor of English, e.g., because they use an English PropBank. Only UCCA is natively fully language independent. Other MRs are being equipped with parallel corpora and node labels (Abzianidze et al., 2020; Van Gysel et al., 2021; Navigli et al., 2022; Giordano and Lopez, 2023). Another trend is to make MRs more expressive. This happens along three avenues: One is extending existing MRFs (as illustrated by AMR extensions for tense or scope modeling, see above), one is to use multi-layer annotation schemes (as exemplified by UCCA or UDS), and one is to employ more complex structures (as DRS does).

However, there is a trade-off between expressivity and the annotation load. To reduce the annotation load, current works aim at crowd-sourcing MRs (e.g., by re-formulating annotation tasks into simple questions (White et al., 2016b)), or improving annotation tools (e.g., with a CodePilot machine-in-the-loop (Cai et al., 2023)), or creating simpler MRFs (Feng et al., 2023).

5.2 Trends in MR Parsing

For humans, producing an MR is an arduous task, particularly for abstract frameworks: a trained annotator needs about 10 minutes to annotate a sentence in AMR Banarescu et al. (2013). Therefore, much research has been dedicated to building automatic parsing systems. Most parsers now use sequence-to-sequence architectures (Ozaki et al., 2020; Samuel and Straka, 2020; van Noord et al., 2018, 2020; Bevilacqua et al., 2021; Zhou et al., 2021). They differ in their learning strategies: graph pre-trainig (Bai et al., 2022; Wang et al., 2023), instruction fine-tuning (Lee et al., 2023), graph information distillation (Vasylenko et al., 2023), or even prompting (Ettinger et al., 2023). Other approaches mix deep learning with classical ideas, using neural representations in transition-based parsing (Astudillo et al., 2020; Zhou et al., 2021), graph-prediction parsing (Lyu and Titov, 2018), or ensembling (Hoang et al., 2021; Lorenzo et al., 2023). Interestingly, despite performance on par with human annotators, recent research suggests that parsing is far from solved (Opitz and Frank, 2022a; Groschwitz et al., 2023).

The evaluation of MRs often revolves around structural graph similarity, measured with metrics such as SMATCH (Cai and Knight, 2013; Opitz, 2023). Computing the SMATCH, however, is NP-complete. Therefore, newer approaches leverage graph-traversal heuristics to evaluate AMR (Song and Gildea, 2019) and DRS (Liu et al., 2020), approximate SMATCH with neural networks (Opitz et al., 2023a), or do quality estimation without a costly reference (Opitz, 2020; Yao and Koller, 2023). There are also efforts towards a more semantic matching of MRs, to take into account, e.g., that a node `cat` is similar to a node `kitten` or a sub-graph `cat :mod young`, with neural networks or graph algorithms used to that end (Opitz et al., 2020, 2021; Shou and Lin, 2023). Recently, Opitz and Frank (2022a) and Groschwitz et al. (2023) showed that SMATCH struggles to detect performance differences between strong parsers.

5.3 Trends in MR Application

We may wonder what is the place of MRFs in a domain dominated by always better-performing LLMs. Some works try to integrate MRs directly into these architectures during the training phase. They leverage the semantic information from MRs to make the models more robust (Zhang et al., 2020b; Cai et al., 2021) or explainable (Opitz and Frank, 2022b). The most popular way for leveraging MR in downstream tasks is to use them as intermediate representations during training and inference. For this, MRs can be linearized and used directly in BERT-like architectures (Ouyang et al., 2021; Xu et al., 2020), or fed into graph neural networks that exploit structure (Song et al., 2019; Xu et al., 2021; Lim et al., 2020; Ribeiro et al., 2022). Other works (Slobodkin et al., 2022) use discourse-level information to perform scene-aware attention, or concatenate sentence and MR embeddings to refine representations (Cai et al., 2022). Again other approaches exploit the graphs directly, in symbolic or neuro-symbolic pipelines. Some works perform MR Parsing and MR-to-text-generation for data augmentation or style transfer (Jangra et al., 2022; Shi et al., 2023; Shou et al., 2022). Others use MRs to do textual inference between pairs of sentences (Bonial et al., 2020; Opitz et al., 2023b), perform splitting for text simplification (Sulem et al., 2018b), or transform MRs into logical formulas to permit symbolic reasoning (Kapanipathi et al., 2021; Chanin and Hunter, 2023).

MRFs are thus being combined fruitfully with LLMs, contributing interpretability, useful intermediate representations, and a bridge towards formal logic.

6 Limitations

Our survey is limited to graph-like meaning representations. While these are indeed the most popular meaning representations these days, there are others that could be discussed in this survey, in particular R. Mooney's ground-breaking works (Mooney, 1996; Wong and Mooney, 2006; Mooney, 2007), or L. Zettlemoyer's work on CCG parsing (Kwiatkowski et al., 2011; Wang et al., 2014; Dasigi et al., 2019), which aim at building Meaning Representations from a corpus, for a target application. The compactness of this survey also prevents us from going more into detail of the parsing techniques. While we do discuss current methods and future trends, parsing itself could merit a survey.

References

Omri Abend and Ari Rappoport. 2013. Universal Conceptual Cognitive Annotation (UCCA). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–238, Sofia, Bulgaria. Association for Computational Linguistics.

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.

Lasha Abzianidze, Johannes Bjerva, Kilian Evang, Hessel Haagsma, Rik van Noord, Pierre Ludmann, Duc-Duy Nguyen, and Johan Bos. 2017. The Parallel Meaning Bank: Towards a multilingual corpus of translations annotated with compositional meaning representations. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 242–247, Valencia, Spain. Association for Computational Linguistics.

Lasha Abzianidze, Johan Bos, and Stephan Oepen. 2020. DRS at MRP 2020: Dressing up discourse representation structures as graphs. In *Proceedings of the CoNLL 2020 Shared Task: Cross-Framework Meaning Representation Parsing*, pages 23–32, Online. Association for Computational Linguistics.

Aashish Arora, Harshitha Malireddi, Daniel Bauer, Asad Sayeed, and Yuval Marton. 2022. Multi-task learning for joint semantic role and proto-role labeling. *arXiv preprint arXiv:2210.07270*.

Nicolas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge.

Ramón Astudillo, Miguel Ballesteros, Tahira Naseem, Austin Blodgett, and Radu Florian. 2020. Transition-based parsing with stack-transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1001–1007, Online. Association for Computational Linguistics.

Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022. Graph pre-training for AMR parsing and generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6001–6015, Dublin, Ireland. Association for Computational Linguistics.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *COLING 1998 Volume 1: The 17th International Conference on Computational Linguistics*.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.

Guntis Barzdins and Didzis Gosko. 2016. RIGA at SemEval-2016 task 8: Impact of Smatch extensions and character-level neural translation on AMR parsing accuracy. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1143–1147, San Diego, California. Association for Computational Linguistics.

Valerio Basile, Johan Bos, Kilian Evang, and Noortje Venhuizen. 2012. Developing a large semantically annotated corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3196–3200, Istanbul, Turkey. European Language Resources Association (ELRA).

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.

Alexandra Birch, Omri Abend, Ondřej Bojar, and Barry Haddow. 2016. HUME: Human UCCA-based evaluation of machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1264–1274, Austin, Texas. Association for Computational Linguistics.

Necva Bölücü and Burcu Can. 2021. Self-attentive constituency parsing for ucca-based semantic parsing. *CoRR*, abs/2110.00621.

Claire Bonial, Stephanie M. Lukin, David Doughty, Steven Hill, and Clare Voss. 2020. 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.

Johan Bos. 2008. Wide-coverage semantic analysis with Boxer. In *Semantics in Text Processing. STEP 2008 Conference Proceedings*, pages 277–286. College Publications.

Johan Bos. 2023. The sequence notation: Catching complex meanings in simple graphs. In *The 15th International Conference on Computational Semantics (IWCS 2023)*.

Chloé Braud, Maximin Coavoux, and Anders Søgaard. 2017. Cross-lingual RST discourse parsing. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 292–304, Valencia, Spain. Association for Computational Linguistics.

Aljoscha Burchardt and Marco Pennacchiotti. 2008. FATE: a FrameNet-annotated corpus for textual entailment. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).

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

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.

Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware? In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2753–2765, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Jon Cai, Shafiuddin Rehan Ahmed, Julia Bonn, Kristin Wright-Bettner, Martha Palmer, and James H. Martin. 2023. CAMRA: Copilot for AMR annotation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 381–388, Singapore. 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.

Yitao Cai, Yue Cao, and Xiaojun Wan. 2021. Revisiting pivot-based paraphrase generation: Language is not the only optional pivot. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4255–4268, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Lynn Carlson, Daniel Marcu, and Mary Ellen Okurovsky. 2001. Building a discourse-tagged corpus in the framework of Rhetorical Structure Theory. In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue*.

Xavier Carreras and Lluís Màrquez. 2004. Introduction to the CoNLL-2004 shared task: Semantic role labeling. In *Proceedings of the Eighth Conference on Computational Natural Language Learning (CoNLL-2004) at HLT-NAACL 2004*, pages 89–97, Boston, Massachusetts, USA. Association for Computational Linguistics.

Xavier Carreras and Lluís Màrquez. 2005. Introduction to the CoNLL-2005 shared task: Semantic role labeling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 152–164, Ann Arbor, Michigan. Association for Computational Linguistics.

Tuhin Chakrabarty, Christopher Hidey, Smaranda Muresan, Kathy McKeown, and Alyssa Hwang. 2019. AMPERSAND: Argument mining for PERSuAsive oNline discussions. 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 2933–2943, Hong Kong, China. Association for Computational Linguistics.

David Chanin and Anthony Hunter. 2023. Neuro-symbolic commonsense social reasoning.

Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A Sag. 2005. Minimal recursion semantics: An introduction. *Research on language and computation*, 3:281–332.

Dipanjan Das and Noah A. Smith. 2011. Semi-supervised frame-semantic parsing for unknown predicates. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1435–1444, Portland, Oregon, USA. Association for Computational Linguistics.

Pradeep Dasigi, Matt Gardner, Shikhar Murty, Luke Zettlemoyer, and Eduard Hovy. 2019. Iterative search for weakly supervised semantic parsing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2669–2680, Minneapolis, Minnesota. Association for Computational Linguistics.

Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal Dependencies. *Computational Linguistics*, 47(2):255–308.

Andrea Di Fabio, Simone Conia, and Roberto Navigli. 2019. VerbAtlas: a novel large-scale verbal semantic resource and its application to semantic role labeling. 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 627–637, Hong Kong, China. Association for Computational Linguistics.

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.

David Dowty. 1991. Thematic proto-roles and argument selection. *Language*, 67(3):547–619.

Timothy Dozat and Christopher D. Manning. 2018. Simpler but more accurate semantic dependency parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 484–490, Melbourne, Australia. 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.

Allyson Ettinger, Jena D Hwang, Valentina Pyatkin, Chandra Bhagavatula, and Yejin Choi. 2023. " you are an expert linguistic annotator": Limits of llms as analyzers of abstract meaning representation. *arXiv preprint arXiv:2310.17793*.

Kilian Evang. 2019. Transition-based DRS parsing using stack-LSTMs. In *Proceedings of the IWCS Shared Task on Semantic Parsing*, Gothenburg, Sweden. Association for Computational Linguistics.

Federico Fancellu, Sorcha Gilroy, Adam Lopez, and Mirella Lapata. 2019. Semantic graph parsing with recurrent neural network DAG grammars. 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 2769–2778, Hong Kong, China. Association for Computational Linguistics.

Hao Fei, Shengqiong Wu, Yafeng Ren, Fei Li, and Donghong Ji. 2021. Better combine them together! integrating syntactic constituency and dependency representations for semantic role labeling. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 549–559, Online. Association for Computational Linguistics.

Lydia Feng, Gregor Williamson, Han He, and Jinho D Choi. 2023. Widely interpretable semantic representation: Frameless meaning representation for broader applicability. *arXiv preprint arXiv:2309.06460*.

Daniel Fernández-González and Carlos Gómez-Rodríguez. 2020. Transition-based semantic dependency parsing with pointer networks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7035–7046, Online. Association for Computational Linguistics.

Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. 2014. A discriminative graph-based parser for the Abstract Meaning Representation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1436, Baltimore, Maryland. Association for Computational Linguistics.

Dan Flickinger, Yi Zhang, and Valia Kordoni. 2012. Deepbank. a dynamically annotated treebank of the wall street journal. In *Proceedings of the 11th International Workshop on Treebanks and Linguistic Theories*, pages 85–96, Lisbon, Portugal.

William Gantt, Lelia Glass, and Aaron Steven White. 2022. Decomposing and recomposing event structure. *Transactions of the Association for Computational Linguistics*, 10:17–34.

Qin Gao and Stephan Vogel. 2011. Corpus expansion for statistical machine translation with semantic role label substitution rules. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 294–298, Portland, Oregon, USA. Association for Computational Linguistics.

Daniel Gildea and Daniel Jurafsky. 2000. Automatic labeling of semantic roles. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 512–520, Hong Kong. Association for Computational Linguistics.

Bastien Giordano and Cédric Lopez. 2023. MR4AP: Meaning representation for application purposes. In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 110–121, Nancy, France. Association for Computational Linguistics.

Venkata Govindarajan, Benjamin Van Durme, and Aaron Steven White. 2019. Decomposing generalization: Models of generic, habitual, and episodic statements. *Transactions of the Association for Computational Linguistics*, 7:501–517.

Jonas Groschwitz, Shay B Cohen, Lucia Donatelli, and Meaghan Fowlie. 2023. Amr parsing is far from solved: Grapes, the granular amr parsing evaluation suite. *arXiv preprint arXiv:2312.03480*.

Jan Hajič, Massimiliano Ciaramita, Richard Johansson, Daisuke Kawahara, Maria Antònia Martí, Lluís

Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Štěpánek, Pavel Straňák, Mihai Surdeanu, Nianwen Xue, and Yi Zhang. 2009. The CoNLL-2009 shared task: Syntactic and semantic dependencies in multiple languages. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009): Shared Task*, pages 1–18, Boulder, Colorado. Association for Computational Linguistics.

Jan Hajič, Eva Hajičová, Jarmila Panevová, Petr Sgall, Ondřej Bojar, Silvie Cinková, Eva Fučíková, Marie Mikulová, Petr Pajas, Jan Popelka, Jiří Semecký, Jana Šindlerová, Jan Štěpánek, Josef Toman, Zdeňka Urešová, and Zdeněk Žabokrtský. 2012. Announcing Prague Czech-English Dependency Treebank 2.0. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3153–3160, Istanbul, Turkey. European Language Resources Association (ELRA).

Silvana Hartmann, Ilia Kuznetsov, Teresa Martin, and Iryna Gurevych. 2017. Out-of-domain FrameNet semantic role labeling. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 471–482, Valencia, Spain. Association for Computational Linguistics.

Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018. Jointly predicting predicates and arguments in neural semantic role labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 364–369, Melbourne, Australia. Association for Computational Linguistics.

Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and what's next. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 473–483, Vancouver, Canada. Association for Computational Linguistics.

Hugo Hernault, Helmut Prendinger, David duVerle, and Mitsuru Ishizuka. 2010. HILDA: A Discourse Parser Using Support Vector Machine Classification. *Dialogue & Discourse*, 1:1–33.

Daniel Hershcovich, Omri Abend, and Ari Rappoport. 2017. A transition-based directed acyclic graph parser for UCCA. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1127–1138, Vancouver, Canada. Association for Computational Linguistics.

Thanh Lam Hoang, Gabriele Picco, Yufang Hou, Young-Suk Lee, Lam Nguyen, Dzung Phan, Vanessa Lopez, and Ramon Fernandez Astudillo. 2021. Ensembling graph predictions for amr parsing. In *Advances in Neural Information Processing Systems*, volume 34, pages 8495–8505. Curran Associates, Inc.

Julia Hockenmaier and Mark Steedman. 2007. CCGbank: A corpus of CCG derivations and dependency structures extracted from the Penn Treebank. *Computational Linguistics*, 33(3):355–396.

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.

Yangfeng Ji and Jacob Eisenstein. 2014. Representation learning for text-level discourse parsing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13–24, Baltimore, Maryland. Association for Computational Linguistics.

Wei Jiang, Zhenghua Li, Yu Zhang, and Min Zhang. 2019. HLT@SUDA at SemEval-2019 task 1: UCCA graph parsing as constituent tree parsing. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 11–15, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Hans Kamp. 1981. A theory of truth and semantic representation. *Formal semantics-the essential readings*, pages 189–222.

Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic*. Springer Dordrecht.

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.*

Karin Kipper, Hoa Trang Dang, and Martha Palmer. 2000. Class-based construction of a verb lexicon. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, page 691–696. AAAI Press.

Naoki Kobayashi, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura, and Masaaki Nagata. 2020. Top-down RST parsing utilizing granularity levels in documents. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(5):8099–8106.

Fajri Koto, Jey Han Lau, and Timothy Baldwin. 2021. Top-down discourse parsing via sequence labelling. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 715–726, Online. Association for Computational Linguistics.

Tom Kwiatkowski, Luke Zettlemoyer, Sharon Goldwater, and Mark Steedman. 2011. Lexical generalization in CCG grammar induction for semantic parsing. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1512–1523, Edinburgh, Scotland, UK. Association for Computational Linguistics.

Young-Suk Lee, Ramón Fernandez Astudillo, Radu Florian, Tahira Naseem, and Salim Roukos. 2023. Amr parsing with instruction fine-tuned pre-trained language models.

Young-Suk Lee, Ramón Fernandez Astudillo, Tahira Naseem, Revanth Gangi Reddy, Radu Florian, and Salim Roukos. 2020. Pushing the limits of AMR parsing with self-learning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3208–3214, Online. Association for Computational Linguistics.

Qi Li, Tianshi Li, and Baobao Chang. 2016. Discourse parsing with attention-based hierarchical neural networks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 362–371, Austin, Texas. Association for Computational Linguistics.

Zuchao Li, Shexia He, Jiaxun Cai, Zhuosheng Zhang, Hai Zhao, Gongshen Liu, Linlin Li, and Luo Si. 2018. A unified syntax-aware framework for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2401–2411, Brussels, Belgium. 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.

Xiang Lin, Shafiq Joty, Prathyusha Jwalapuram, and M Saiful Bari. 2019. A unified linear-time framework for sentence-level discourse parsing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4200, 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.

Jiangming Liu, Shay B. Cohen, and Mirella Lapata. 2020. Dscorer: A fast evaluation metric for discourse representation structure parsing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4547–4554, Online. Association for Computational Linguistics.

ACM Lorenzo, Pere Lluís Huguet Cabot, and Roberto Navigli. 2023. AMRs assemble! learning to ensemble with autoregressive models for AMR parsing. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1595–1605, Toronto, Canada. Association for Computational Linguistics.

Chunchuan Lyu and Ivan Titov. 2018. AMR parsing as graph prediction with latent alignment. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 397–407, Melbourne, Australia. Association for Computational Linguistics.

William C. Mann and Sandra A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text - Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281.

George A. Miller. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11):39–41.

Jelena Mitrović, Cliff O'Reilly, Miljana Mladenović, and Siegfried Handschuh. 2017. Ontological representations of rhetorical figures for argument mining. *Argument & Computation*, 8(3):267–287.

Yusuke Miyao. 2006. *From linguistic theory to syntactic analysis. Corpus-oriented grammar development and feature forest model*. Ph.D. thesis, University of Tokyo, Tokyo, Japan.

Muhidin Mohamed and Mourad Oussalah. 2019. Srl-esa-textsum: A text summarization approach based on semantic role labeling and explicit semantic analysis. *Information Processing & Management*, 56(4):1356–1372.

Raymond J. Mooney. 1996. Inductive logic programming for natural language processing. In Stephen Muggleton, editor, *Inductive Logic Programming: Selected papers from the 6th International Workshop*, pages 3–22. Springer Verlag, Berlin.

Raymond J. Mooney. 2007. Learning for semantic parsing. In *Computational Linguistics and Intelligent Text Processing: Proceedings of the 8th International Conference (CICLing 2007)*, pages 311–324, Mexico City, Mexico. Springer: Berlin, Germany. Invited paper.

Mathieu Morey, Philippe Muller, and Nicholas Asher. 2017. How much progress have we made on RST discourse parsing? a replication study of recent results on the RST-DT. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1319–1324, Copenhagen, Denmark. Association for Computational Linguistics.

Tahira Naseem, Austin Blodgett, Sadhana Kumaravel, Tim O'Gorman, Young-Suk Lee, Jeffrey Flanigan, Ramón Fernandez Astudillo, Radu Florian, Salim Roukos, and Nathan Schneider. 2021. Docamr: Multi-sentence AMR representation and evaluation. *CoRR*, cs.CL/2112.08513.

Roberto Navigli, Michele Bevilacqua, Simone Conia, Dario Montagnini, and Francesco Cecconi. 2021. Ten years of BabelNet: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4559–4567. International Joint Conferences on Artificial Intelligence Organization. Survey Track.

Roberto Navigli, Rexhina Blloshmi, and Abelardo Carlos Martínez Lorenzo. 2022. BabelNet Meaning Representation: A fully semantic formalism to overcome language barriers. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(11):12274–12279.

Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty, and Xiaoli Li. 2021. RST parsing from scratch. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1613–1625, Online. Association for Computational Linguistics.

Stephan Oepen, Omri Abend, Jan Hajic, Daniel Hershcovich, Marco Kuhlmann, Tim O'Gorman, Nianwen Xue, Jayeol Chun, Milan Straka, and Zdenka Uresova. 2019. MRP 2019: Cross-framework meaning representation parsing. In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning*, pages 1–27, Hong Kong. Association for Computational Linguistics.

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinková, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Zdeňka Urešová. 2016. Towards comparability of linguistic graph Banks for semantic parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3991–3995, Portorož, Slovenia. European Language Resources Association (ELRA).

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinková, Dan Flickinger, Jan Hajič, and Zdeňka Urešová. 2015. SemEval 2015 task 18: Broad-coverage semantic dependency parsing. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 915–926, Denver, Colorado. Association for Computational Linguistics.

Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Yi Zhang. 2014. SemEval 2014 task 8: Broad-coverage semantic dependency parsing. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 63–72, Dublin, Ireland. Association for Computational Linguistics.

Stephan Oepen and Jan Tore Lønning. 2006. Discriminant-based MRS banking. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).

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.

Juri Opitz. 2020. AMR quality rating with a lightweight CNN. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 235–247, Suzhou, China. Association for Computational Linguistics.

Juri Opitz. 2023. SMATCH++: Standardized and extended evaluation of semantic graphs. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1595–1607, Dubrovnik, Croatia. Association for Computational Linguistics.

Juri Opitz, Angel Daza, and Anette Frank. 2021. 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. 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 Meier, and Anette Frank. 2023a. SMARAGD: Learning SMatch for accurate and rapid approximate graph distance. In *Proceedings of the 15th International Conference on Computational Semantics*, pages 267–274, Nancy, France. 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. 2023b. 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.

Siru Ouyang, Zhuosheng Zhang, and Hai Zhao. 2021. Dialogue graph modeling for conversational machine reading. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3158–3169, Online. Association for Computational Linguistics.

Hiroaki Ozaki, Gaku Morio, Yuta Koreeda, Terufumi Morishita, and Toshinori Miyoshi. 2020. Hitachi at MRP 2020: Text-to-graph-notation transducer. In *Proceedings of the CoNLL 2020 Shared Task: Cross-Framework Meaning Representation Parsing*, pages 40–52, Online. Association for Computational Linguistics.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.

Siyana Pavlova, Maxime Amblard, and Bruno Guillaume. 2023. Structural and Global Features for Comparing Semantic Representation Formalisms. In *The 4th International Workshop on Designing Meaning Representation*, Nancy, France.

Andreas Peldszus and Manfred Stede. 2013. From argument diagrams to argumentation mining in texts: A survey. *International Journal of Cognitive Informatics and Natural Intelligence (IJCINI)*, 7(1):1–31.

Xiaochang Peng, Linfeng Song, Daniel Gildea, and Giorgio Satta. 2018. Sequence-to-sequence models for cache transition systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1842–1852, Melbourne, Australia. Association for Computational Linguistics.

Lekshmi R Pillai, Veena G., and Deepa Gupta. 2018. A combined approach using semantic role labelling and word sense disambiguation for question generation and answer extraction. In *2018 Second International Conference on Advances in Electronics, Computers and Communications (ICAECC)*, pages 1–6.

Sameer Pradhan, Wayne Ward, Kadri Hacioglu, James Martin, and Daniel Jurafsky. 2005. Semantic role labeling using different syntactic views. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 581–588, Ann Arbor, Michigan. Association for Computational Linguistics.

Vasin Punyakanok, Dan Roth, and Wen-tau Yih. 2008. The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics*, 34(2):257–287.

James Pustejovsky, Ken Lai, and Nianwen Xue. 2019. Modeling quantification and scope in Abstract Meaning Representations. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 28–33, Florence, Italy. Association for Computational Linguistics.

Drew Reisinger, Rachel Rudinger, Francis Ferraro, Craig Harman, Kyle Rawlins, and Benjamin Van Durme. 2015. Semantic proto-roles. *Transactions of the Association for Computational Linguistics*, 3:475–488.

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.

Rachel Rudinger, Adam Teichert, Ryan Culkin, Sheng Zhang, and Benjamin Van Durme. 2018. Neural-Davidsonian semantic proto-role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 944–955, Brussels, Belgium. Association for Computational Linguistics.

David Samuel and Milan Straka. 2020. ÚFAL at MRP 2020: Permutation-invariant semantic parsing in PERIN. In *Proceedings of the CoNLL 2020 Shared Task: Cross-Framework Meaning Representation Parsing*, pages 53–64, Online. Association for Computational Linguistics.

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.

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.

Aviv Slobodkin, Leshem Choshen, and Omri Abend. 2022. Semantics-aware attention improves neural machine translation. In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, pages 28–43, Seattle, Washington. Association for Computational Linguistics.

Linfeng Song and Daniel Gildea. 2019. SemBleu: A robust metric for AMR parsing evaluation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4547–4552, Florence, Italy. Association for Computational Linguistics.

Linfeng Song, Daniel Gildea, Yue Zhang, Zhiguo Wang, and Jinsong Su. 2019. Semantic Neural Machine Translation Using AMR. *Transactions of the Association for Computational Linguistics*, 7:19–31.

Radu Soricut and Daniel Marcu. 2003. Sentence level discourse parsing using syntactic and lexical information. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 228–235.

Elizabeth Spaulding, Gary Kazantsev, and Mark Dredze. 2023. Joint end-to-end semantic proto-role labeling. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 723–736, Toronto, Canada. Association for Computational Linguistics.

Elias Stengel-Eskin, Kenton Murray, Sheng Zhang, Aaron Steven White, and Benjamin Van Durme. 2021. Joint universal syntactic and semantic parsing. *Transactions of the Association for Computational Linguistics*, 9:756–773.

Elias Stengel-Eskin, Aaron Steven White, Sheng Zhang, and Benjamin Van Durme. 2020. Universal decompositional semantic parsing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8427–8439, Online. Association for Computational Linguistics.

Elior Sulem, Omri Abend, and Ari Rappoport. 2018a. Semantic structural evaluation for text simplification. 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 685–696, New Orleans, Louisiana. Association for Computational Linguistics.

Elior Sulem, Omri Abend, and Ari Rappoport. 2018b. Simple and effective text simplification using semantic and neural methods. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 162–173, Melbourne, Australia. Association for Computational Linguistics.

Mihai Surdeanu, Richard Johansson, Adam Meyers, Lluís Màrquez, and Joakim Nivre. 2008. The CoNLL 2008 shared task on joint parsing of syntactic and semantic dependencies. In *CoNLL 2008: Proceedings of the Twelfth Conference on Computational Natural Language Learning*, pages 159–177, Manchester, England. Coling 2008 Organizing Committee.

Zhixing Tan, Mingxuan Wang, Jun Xie, Yidong Chen, and Xiaodong Shi. 2018. Deep semantic role labeling with self-attention. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1):4929–4936.

Jens E. L. 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*, 35(3):343–360.

Rik van Noord. 2019. Neural boxer at the IWCS shared task on DRS parsing. In *Proceedings of the IWCS Shared Task on Semantic Parsing*, Gothenburg, Sweden. Association for Computational Linguistics.

Rik van Noord, Lasha Abzianidze, Antonio Toral, and Johan Bos. 2018. Exploring neural methods for parsing discourse representation structures. *Transactions of the Association for Computational Linguistics*, 6:619–633.

Rik van Noord, Antonio Toral, and Johan Bos. 2020. Character-level representations improve DRS-based semantic parsing even in the age of BERT. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4587–4603, Online. Association for Computational Linguistics.

Siddharth Vashishtha, Benjamin Van Durme, and Aaron Steven White. 2019. Fine-grained temporal relation extraction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2906–2919, Florence, Italy. Association for Computational Linguistics.

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.

David Vilares and Carlos Gómez-Rodríguez. 2018. A transition-based algorithm for unrestricted AMR parsing. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 142–149, New Orleans, Louisiana. Association for Computational Linguistics.

Adrienne Wang, Tom Kwiatkowski, and Luke Zettlemoyer. 2014. Morpho-syntactic lexical generalization for CCG semantic parsing. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1284–1295, Doha, Qatar. Association for Computational Linguistics.

Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015. A transition-based algorithm for AMR parsing. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 366–375, Denver, Colorado. Association for Computational Linguistics.

Chunliu Wang, Huiyuan Lai, Malvina Nissim, and Johan Bos. 2023. Pre-trained language-meaning models for multilingual parsing and generation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5586–5600, Toronto, Canada. Association for Computational Linguistics.

Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. Automated Concatenation of Embeddings for Structured Prediction. In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. Association for Computational Linguistics.

Yizhong Wang, Sujian Li, and Houfeng Wang. 2017. A two-stage parsing method for text-level discourse analysis. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 184–188, Vancouver, Canada. Association for Computational Linguistics.

Yizhong Wang, Sujian Li, and Jingfeng Yang. 2018. Toward fast and accurate neural discourse segmentation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 962–967, Brussels, Belgium. Association for Computational Linguistics.

Keenon Werling, Gabor Angeli, and Christopher D. Manning. 2015. Robust subgraph generation improves Abstract Meaning Representation parsing. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 982–991, Beijing, China. Association for Computational Linguistics.

Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016a. Universal decompositional semantics on Universal Dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723, Austin, Texas. Association for Computational Linguistics.

Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016b. Universal decompositional semantics on Universal Dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723, Austin, Texas. Association for Computational Linguistics.

Aaron Steven White, Elias Stengel-Eskin, Siddharth Vashishtha, Venkata Subrahmanyan Govindarajan, Dee Ann Reisinger, Tim Vieira, Keisuke Sakaguchi, Sheng Zhang, Francis Ferraro, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2020. The universal decompositional semantics dataset and decomp toolkit. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 5698–5707, Marseille, France. European Language Resources Association.

Yuk Wah Wong and Raymond Mooney. 2006. Learning for semantic parsing with statistical machine translation. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 439–446, New York City, USA. Association for Computational Linguistics.

Jiacheng Xu, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Discourse-aware neural extractive text summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5021–5031, Online. Association for Computational Linguistics.

Weiwen Xu, Huihui Zhang, Deng Cai, and Wai Lam. 2021. Dynamic semantic graph construction and reasoning for explainable multi-hop science question answering. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1044–1056, Online. Association for Computational Linguistics.

Yuekun Yao and Alexander Koller. 2023. Predicting generalization performance with correctness discriminators. *arXiv preprint arXiv:2311.09422*.

Nan Yu, Meishan Zhang, and Guohong Fu. 2018. Transition-based neural RST parsing with implicit syntax features. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 559–570, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Longyin Zhang, Yuqing Xing, Fang Kong, Peifeng Li, and Guodong Zhou. 2020a. A top-down neural architecture towards text-level parsing of discourse rhetorical structure. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6386–6395, Online. Association for Computational Linguistics.

Sheng Zhang, Xutai Ma, Rachel Rudinger, Kevin Duh, and Benjamin Van Durme. 2018. Cross-lingual decompositional semantic parsing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1664–1675, Brussels, Belgium. Association for Computational Linguistics.

Yu Zhang, Qingrong Xia, Shilin Zhou, Yong Jiang, Zhenghua Li, Guohong Fu, and Min Zhang. 2021. Semantic role labeling as dependency parsing: Exploring latent tree structures inside arguments. *CoRR*, abs/2110.06865.

Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020b. Semantics-aware bert for language understanding. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9628–9635.

Wanjun Zhong, Jingjing Xu, Duyu Tang, Zenan Xu, Nan Duan, Ming Zhou, Jiahai Wang, and Jian Yin. 2020. Reasoning over semantic-level graph for fact checking. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6170–6180, Online. Association for Computational Linguistics.

Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo, and Radu Florian. 2021. AMR parsing with action-pointer transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5585–5598, Online. Association for Computational Linguistics.

Jie Zhou and Wei Xu. 2015. End-to-end learning of semantic role labeling using recurrent neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1127–1137, Beijing, China. Association for Computational Linguistics.