Syntax is the Key to Semantics: Universal Dependencies and Abstract Meaning Representation

Johannes Heinecke
Orange Research, 22300 Lannion, France
johannes.heinecke@orange.com

Abstract

This paper presents a new Abstract Meaning Representation (AMR) dataset using sentences which are already translated in 21 languages and annotated in dependency syntax (Parallel Universal Dependencies, PUD). We annotated the English version of the 1000 sentences available in the PUD dataset. Since PUD provides syntactic annotations, this new AMR dataset (ParallelUD-AMR) allows comparisons of syntactic phenomena and their corresponding semantics. We finally provide a first analysis of parallels between syntactic and semantic structures.

1. Introduction

Syntax is an important feature of all natural languages: word order, agreement, relations between words (adjacent or not), hierarchies etc. In most cases changing the syntactic construction of a sentence implies a change of meaning.

Parallel UD helps to compare and see differences in syntax and typology of languages. In order to have directly comparable corpora in syntax and semantics, we created a new Abstract Meaning Representation (AMR) dataset named ParallelUD-AMR[1], which is based on the sentences UD_English-PUD treebank[2]. We annotated the 1000 sentences following the AMR guidelines[3]. The goal of this new dataset is to show parallels of Semantic and Syntactic representations in different languages. In the following we give a short overview on related work (section 2), present existing UD and AMR data (3) and the annotation process (4). Section 4 also details some inconsistencies in the PUD translations which in consequence mean that the AMR graph calqued on the English version of the sentence does not correspond exactly to some of the translated sentences. Section 5 presents a first idea of how this dataset can be used for comparison. We conclude with some ideas of how this dataset could also be used (section 6.

2. Background

Even though Universal Dependencies and Abstract Meaning Representation annotate different linguistic layers (dependency syntax in the former case, argument semantics in the latter) there is an interest to combine them, since at least a part of the semantic relations which are represented by the AMR graphs are a consequence of the syntactic construction of the sentence. For instance the subject of the English verb *to see* in the Active Voice encodes the agent of the *seeing*-event (coded `:ARG0` in AMR), whereas the subject of the same verb in the Passive Voice encodes the patient (coded `:ARG1` in AMR), cf. Schuster and Manning (2016) for a comparison of UD data and AMR. Alignment between AMR and dependency syntax has already been discussed in Szubert et al. (2018) who discuss how to establish relations between syntactic dependencies and semantic (AMR) relations. We think our new corpus can help to advance in this direction. Parallel Universal Dependencies data (PUD, see below in section 3.1) has been used for a shared task a while ago (Mille et al., 2018). Instead of creating semantic graph; participants were asked to add word order and inflected forms from incomplete dependency trees. Even though this task does not involve AMR or a different kind of semantics it shows the interest which lies in multilingual parallel data. Other work exists to at least bootstrap or even create AMR graphs by using a syntactic representations of sentence (Pust et al., 2015), frequently taken from UD (Han and Pavlova, 2019). Further approaches use UD data to bootstrap Uniform Meaning Representation (UMR, (Van Gysel et al., 2021)), a (multilingual) evolution of AMR (Gamba et al., 2025). Wang et al. (2017) published an approach of a parallel use of dependency syntax trees and AMR for information extraction, showing the interest in corpora annotated in both frameworks. A complementary approach is presented by Wein and Schneider (2024). They discuss differences and similarities in parallel (multilingual) AMRs also taken from translations but using concepts and relations drawn from the sentence languages (and not En-

[1] The dataset will be available at https://github.com/orange-Opensource/ParallelUD-AMR.

[2] https://github.com/UniversalDependencies/UD_English-PUD/

[3] https://github.com/amrisi/amr-guidelines/blob/master/amr.md

Proceedings of the Ninth Workshop on Universal Dependencies (UDW 2026), pages 295–304
May 16, 2026. ©ELRA Language Resources Association (ELRA), 2026

glish as in our corpus). However, there is no formal syntax in their approach.

Two early AMR-Parsers were initially based on transition-parsing developped for the parsing of dependency syntax: C-AMR (Wang et al., 2015)[4] and JAMR[5]. Both show the interest of regarding syntax and semantics at the same time.

3. Data

3.1. Universal Dependencies (UD)

Universal Dependencies is a large community activity to annotate as many languages as possible in dependency syntax, using the same annotation guidelines (Nivre et al., 2020; Marneffe et al., 2021). The latest release (2.17 of November 15, 2025[6]), contains over 200 treebanks for more than 150 languages. If a treebank is large enough, a split in training/dev/test is available to train dependency parsers like UDpipe (Straka and Straková, 2017; Straka, 2018) or Stanza (Qi et al., 2020).

3.2. Parallel Universal Dependencies (PUD)

PUD is an initiative within the UD project to annotate 1000 sentences translated into currently 21 languages of different language families (Arabic, Chinese, Czech, English, Finnish, French, Galician, German, Hindi, Indonesian, Icelandic, Italian, Japanese, Korean, Polish, Portuguese, Russian, Spanish, Swedish, Thai and Turkish). Half of the sentences is taken from news, the other half stems from Wikipedia. The 1000 sentences appear in the same order and with the same sentence identification in the respective files. 750 of the original sentences were English. The rest was taken from German, French, Italian or Spanish resources[7] and translated via English into all other languages by professional translators. The translations were then annotated according the UD v2 guidelines and are available via the UD website[8]. We used release 2.16 for this work. There are other datasets available with different translations[9], PUD however is the largest dataset (both in number of sentences and number of different languages).

[4]https://github.com/c-amr/camr

[5]https://github.com/jflanigan/jamr

[6]We use release 2.16 in this paper (Zeman, 2025)

[7]The first two digits of the sentence ID indicate the original language: English (01), German (02), French (03), Italian (04) or Spanish (05).

[8]https://universaldependencies.org

[9]https://universaldependencies.org/contributing/parallel.html

3.3. Abstract Meaning Representation (AMR)

Abstract Meaning Representation (Banarescu et al., 2013) provides a framework to model the meaning of a sentence, notably actions, events or states and their participants. AMR relies heavily on (verbal) concepts defined in PropBank (Kingsbury and Palmer, 2002; Palmer et al., 2005)[10], e.g. the concept `cost-01` in figure 1 (serialised PENMAN format (Kasper, 1989)) refers to PropBank's sense (or "roleset") `-01` of the verb "to cost". Instances are indicated by a following "/", e.g., `c` being an instance of the concept `cost-01`. The names of the variables do not have any other semantics than being distinct. Relations are indicated by an initial colon (e.g. `:ARG1`, `:time`). Literals (strings and numbers) lack a preceding instance and "/" (c.f. *"NoMa"* and *2004* in the example in figure 1). Alignments with tokens of the sentence are optional in PENMAN and are marked by adding `~e.n` after the concept, relation or literal. *n* aligns the instance/relation/literal with the *n*th token (0-indexed) or the sentence.

```
# 0  1           2 3  4    5 6     7       8
# By comparison , it cost $ 103.7 million to
# 9     10  11    12      13    14      15
# build the NoMa infill Metro station ,
# 16    17     18 19   20
# which opened in 2004 .
(c2 / cost-01~e.4
   :ARG1 (b / build-01~e.9
      :ARG1 (s / station~e.14
         :ARG1-of (o / open-01~e.17
            :time~e.18 (d / date-entity
               :year 2004~e.19))
            :mod (i / infill~e.12)
            :mod (s1 / subway~e.13)
            :name (n / name
               :op1 "NoMa"~e.11)))
   :ARG2 (m / monetary-quantity
      :quant 103700000~e.6,7
      :unit (d2 / dollar~e.5))
   :ARG2-of (c / compare-01~e.1))
```

Figure 1: AMR graph for "By comparison, it cost $103.7 million to build the NoMa infill Metro station, which opened in 2004." in PENMAN format including alignments with tokens (sentence taken from UD-English-PUD id n01005023; our AMR annotation)

The largest available corpus used to train models capable of parsing sentences from natural languages into AMR graphs, called AMR 3.0, LDC2020T02 (Knight et al., 2020)[11], is provided by the Linguistic Data Consortium (LDC). This corpus is composed of nearly 59 000 sentences and

[10]https://propbank.github.io/

[11]https://catalog.ldc.upenn.edu/LDC2020T02

corresponding AMR graphs. The data contains discussions from forums (partly technical), news reels, translations to English of Chinese news broadcasts, along with a part originating from English Wikipedia pages and Aesop's fables (see LDC2020T02 documentation).

Even though AMR has been designed originally for English, it has also been employed successfully for other languages. A translation of the test dataset of AMR-2.0 is available at LDC (Damonte and Cohen, 2020)[12].

In order annotate languages other than English, there are two approaches: Using concepts in the language of the sentences (by using a corresponding version of PropBank, if existing, e.g. Akbik et al. (2015)[13], or keeping concepts in English (Heinecke and Shimorina, 2022).

An extension of AMR, Uniform Meaning Representation (UMR) has been designed for multilingual usage by keeping the concepts of the graph in the same language as the corresponding sentences (Van Gysel et al., 2021). UMR includes all of the information from AMR, and additional information about the position of the sentence in a document, temporal layering and coreferences between entities of different sentences and alignments of tokens of the sentence with instances of the UMR graph.

Whereas UMR follows the first approach we chose the latter for the annotation of the PUD data, since the available translations of the PUD sentences are supposed to have the same meaning, so a single AMR graph for each sentences and its translations is needed. Since PropBank is available only in English and equivalents in other languages are less complete or not available, AMR graphs based on English seems to be the best solution.

4. Annotation

Out of the 1000 sentences taken from UD_English-PUD the first part of 500 sentences were annotated by 4 annotators, using the AMR annotation tool *metAMoRphosED* (Heinecke, 2023), the second part of 500 sentences was annotated by 3 annotators. The annotators did not have to write the graph entirely. We used two different AMR predictors to provide different preannotations to the annotators. For this we used AMRlib[14] trained on AMR 3.0 and AMR-transition parser (Drozdov et al., 2022; Lee et al., 2022)[15]. The annotators task was to verify and if needed to modify the generated graph with respect to the AMR guidelines.

During and after the annotation we checked the AMR graph to ensure some basic formal validity (either provided by the annotation tool *metAMoRphosED* or by additional scripts). In particular we checked whether

- concepts with a numerical suffix (like `see-01` or `have-org-rel-91` actually exist in PropBank[16] and that eventually outgoing `:ARGn`-relations are defined for the given text.
- there are no duplicate `:ARGn` between two instances `(e/eat-01 :ARG1 (...) :ARG1 (....))`
- all relations other than `:ARGn` are valid AMR relations
- named entities in the sentence are correctly represented in the Graph (including the fact that demonyms may appear as adjectives in the text (e.g. *American* or *German*. They are nevertheless represented by `(c/country :name (n/name :op1 "United" :op1 "States" :op2 "of" :op4 "America"))` or `((c/country :name (n/name :op1 "Germany"))`
- dates in the text correspond to `date-entity` instances in the graph: E.g. *...June 2002...* is represented as `(d/date-entity :month 6 :year 2002)`
- numerical values in the text correspond to targets of `:quant` or `:value` relations: E.g. *two million dollars* is represented by `(m/monetary-quantity :quant 2000000 :unit (d/dollar))`

In terms of Smatch (Cai and Knight, 2013) the inter-annotator agreement (IAA) (Wein and Schneider, 2024, p. 421) was 92.4% for the former and 93.1% for the latter half using average pairwise agreement. The upper graph in Fig. 2 presents the averaged Smatch between all 4 annotators for the first 500 sentences. Since the Smatch metric depends on the graph size[17] we also show the averaged number of differences between all annotators. In spite of some outliers, in general annotators agreed well. Fig. 3 show the same information for the second 500 sentences (3 annotators).

In cases where not all 4 (or 3) annotators agreed, we manually chose one of the graphs. We also compared the final annotations of each annotator with the original pre-annotated files provided to each of the annotators. IAA varies between 93.1% (2.6 differences per graph in average) and

[12] https://catalog.ldc.upenn.edu/LDC2020T07
[13] https://github.com/UniversalPropositions/UP-1.0
[14] https://github.com/bjascob/amrlib
[15] https://github.com/IBM/transition-amr-parser
[16] https://github.com/propbank/propbank-frames/
[17] A single error in a large graph decreases the Smatch score only slightly, whereas a single error in a small graph decreases the score drastically.

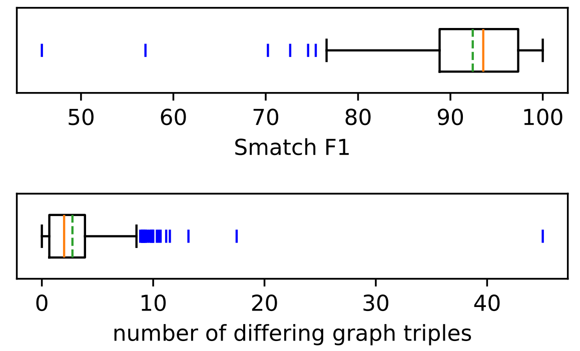

Figure 2: Boxplot average of pairwise IAA for first part of the corpus. Top: Smatch F1: Median (solid orange line): 93.6%, mean (dashed green line): 92.4%, std. deviation: 6.52 Bottom: number of differing triples between graphs: Median: 2.0, mean: 2.75, std. deviation: 3.18

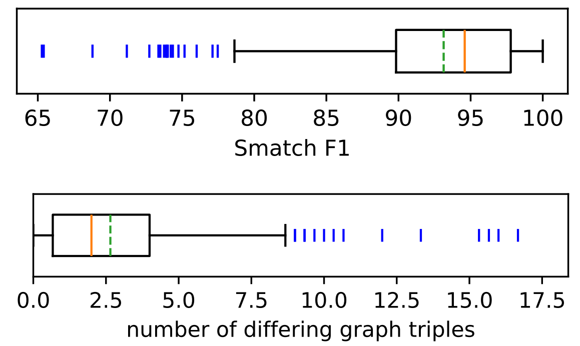

Figure 3: Boxplot average of pairwise IAA for second part of the corpus. Top: Smatch F1: Median (solid orange line): 94.6%, mean (dashed green line): 93.1%, std. deviation: 6.44 Bottom: number of differing triples between graphs: Median: 2.0, mean: 2.65, std. deviation: 2.70

96.9% (0.62 differences per graph). This means that the annotators did not completely rewrote the automatic generated graph, but In a final step we added alignments between words of the sentence and instances, relations and literals of the corresponding AMR graph (following the alignment guidelines of AMR 3.0).

We chose the English PUD sentences as the basis for the annotation for two reasons: 1) 750 of the sentences were originally already in English; 2) AMR was designed for English, even though it is used to annotate sentences in other languages

4.1. Inconsistencies

Even though the translations into the other languages of the PUD treebanks were produced by professional translators, there are cases where a more literal translation would have facilitated the alignment between the AMR graph and the non-English sentence. These minor differences mean that the AMR graph does not correspond as well to the non-English sentence than to the English sentence. For instance (the first language in these examples is the original language of the sentence; list not exhaustive):

translation of Named Entities: There are cases where a named entity (NE) in English was translated by a different expression or the function of the NE:

- en: *on Capitol Hill* (id: n01001013) →
  - fr: *sur **le Capitol***;
  - fi: *Yhdysvaltain **kongressissa*** “in the US congress”
- en: *the **GOP** nominee* (id: n01002017) →
  - de: *der Nominierte der **Republikanischen Partei***,
  - ru: *кандидат **Республиканской партии***,
  - tr: ***Cumhuriyetçi Parti** Başkan adayı Trump* “the Republican Party presidential candidate Trump” the Turkish translation also mentions the name “Trump”, whereas in the other languages only a “nominee” or “candidate” is mentioned
- en: *Zettel’s Traum* (“the dream of Zettel”, here “Zettel” being a person name, but in German the noun “Zettel” also means small bit of paper for notes) (id: n01066045) →
  - es: *El sueño de **la ficha*** “the dream of the card/paper”);
  - is: *Draumi **Spóla*** “the dream of the tape”

incorrect translation of Named Entities:

- en: *It’s the election, of course, not ”Game of Thrones.”* (id: n01132006) →
  - pt: *São as eleições, claro, não a ”**Guerra** dos Tronos.”*

adding details to Named Entities: Titles and first names appear in translations but were absent in the original (English) sentence:

- en: *Mr Panvalkar* (id: n01010042) →
  - de: ***Prasad** Panvalker*
- en: *Lamarr accompanied Mandl to business meetings* (id: w01101020) →

- fr: ***Hedy** Lamarr accompagnait **Friedrich** Mandl aux réunions d'affaires*

- en: *De Gaulle had instructed* →
  - cs: ***Charles** de Gaulle vydal pokyny*
- en: *until after Smith's death.* (id: 01139043) →
  - fr: *jusqu'après la mort de **E.E.** Smith*
- en: *Pebe moved the family to Nashville, Tennessee* (id: w01147010) →
  - fr: *En 1991, Pebe **Seber** déménage avec sa famille à Nashville (Tennessee)*

adding/changing types of Named Entities:

- en: *Wright was born in Poole, Dorset...* (id: w01149002) →
  - ar: ولد رايت في **مدينة** بول "Wright was born in the **city** of Poole"
- fr: *Le maire de Saint-Gaudens (Haute-Garonne)* (id: n03004003) →
  - de: *Bürgermeister von Saint-Gaudens (**im Département** Haute-Garonne)*,
  - es: *el alcalde de Saint-Gaudens (**en la región** Haute-Garonne)*,
  - ar: عمدة سان - غودانس (**في إقليم** العليا غارون) "the **region** Haute-Garonne",
  - ru: *мэр города Сан-Годан (**в регионе** Верхняя Гаронна)*
- de: *Larry Vance etwa sagte dem SPIEGEL* (id: n02066010) →
  - ru: *Ларри Вэнс заявил **журналу** SPIEGEL*

removing type from Named Entities: The opposite of the case before:

- es: *federalización de **la Ciudad de** Buenos Aires* (id: w05008107) →
  - cs: *federalizaci Buenos Aires*,
  - ru: *федерализацию Буэнос-Айреса*
- de: ***Städte** Köln und Bremen* (id: w02018030) →
  - ru: *Кёльна и Бремена*

added preposition in Named Entities:

- en: CTV Montreal →
  - ru: *CTV **в** Монреале* (id: w01141025)

"free" translation: The translation is adapted to the cultural context and the way things are expressed in the target language and is not a literal translation

- en: ***large** bank account* (id: n01010042) →
  - de: ***wohlgefülltes** Bankkonto* "well filled bank account",
  - fr: *compte en banque **imposant*** "impressing bank account",
  - ru: ***крупным суммам** на банковском счету* "**large sums** in [C.'s] bank account"
- en: *More than 330 crew are **onboard** the ship.* (id: n01070016) →
  - fr: *L'équipage du navire **comporte** plus de 330 membres.*
- it: *Italia abbia **meno** chilometri di metropolitana* (id: n04002020) →
  - de: *dass Italiens U-Bahnlinien **kürzer** sind*,
  - tr: *İtalya'nın **daha** kısa bir metro hattı olduğunu*
- de: *steigenden **Staats**ausgaben* (id: n02078011) →
  - en: *increasing **federal** expenditures*, es:*aumento del gasto del **gobierno federal***

incorrect translation: In contrast to the former, here a term has been incorrectly translated

- de: *Kloster Eibingen* (id: w02015087 also w02004008, w02015084 and w02015086) →
  - en: *cloister Eibingen*, fr:*cloître d'Eibingen* (instead of *monastary/monastère Eibingen*),

"incomplete" translation: Words or expression which are not translated

- en: *meanwhile* not translated in German (id: n01101017)
- de: *dabei* not translated in any otther language (id: w02019077)

"overcomplete" translation: Additions absent in the English source or implicits made explicit

- en: *Is series two working so far?* (id: n01121051) →

- fr: *Ça marche **pour toi** jusqu'à présent, la deuxième saison ?*

- en: *For those who follow social media transitions on Capitol Hill, this will be a little different.* (id: n01001013) →
  - tr: ***Bu durum** Capitol Hill hakkındaki sosyal medya hareketlerini takip edenler için biraz daha farklı olacak.* "**this situation** will be a little different"; cf. Fig. 4 and 5.

"tautological" translation:

- en: *The **northern** portion of the Caucasus is known as the **Cis**caucasus and the southern portion as the Transcaucasus.* (id: w01014002) →
  - ru: ***Северная** часть Кавказа известна как **Северный Кавказ**, а южная часть — как Закавказье.*

typographical errors:

- fr: *Au VIIIe siècle **av[ant]** J.-C.* "before Christ") (id: w03005012, w03005014) →
  - de: *Im 8. Jahrhundert **n[ach]** Chr.* "after Christ"
- fr: *En pluviôse an II* (id: w03008029)→
  - de: *Im Monat **Pluvoise** des Jahres **III*** (in several languages the French original *an II* "year two" is translated with "year III";

5. Comparison

As mentioned in the introduction, the motivation to create the ParallelUD-AMR dataset was to be able to identify parallels (and differences) between the dependency trees and the semantic Graph (in AMR). Whereas the syntactic annotation is anchored to the words (tokens) of the sentence, the AMR graph (at least in its basic form) is not anchored to its corresponding sentence which makes it difficult for any automatic alignment of syntactic and semantic relations. However the AMR formalism provides the alignment of tokens of the sentence with instances, literals or relations of the AMR graph (a part of the AMR 3.0 dataset comes with alignments).

We started to compare the AMR relations (between instances of concepts or literal) with the dependency relations (between tokens of the sentence). In order to do so, we not only needed alignments between the tokens of the English sentence and the AMR graph (done during the AMR annotation) but also alignments for the languages we wanted to compare with AMR. Figures 4 and 5 present these alignments for English and Turkish.

While we did the alignments for English manually, for other languages (in our case French, Arabic and Turkish) we used an LLM (GPT-4.1-mini) to provide alignments. We validated the entire French corpus manually but only some sentences of the Arabic and Turkish corpus. We used the following prompt in Fig. 7 with two examples (sentence and the corresponding AMR graph with alignments).

System prompt: *you are a specialist in Abstract Meaning Representation Semantics.*
User prompt:
Please give me a list of correspondences between the words of the French sentence and an AMR-concept of the semantic graph (use a json dictionary for the output):
Two examples:
- sentence: Le président préfère de ne pas organiser des éléctions anticipées .
- semantic graph:

```
( p / prefer-01
 :ARG0 ( p2 / person
   :ARG0-of ( h / have-org-role-91
     :ARG2 ( p3 / president)))
 :ARG1 ( h2 / hold-04
   :polarity -
   :ARG0 p2
   :ARG1 ( e / elect-01
     :time ( e2 / early))))
```

- output:

```
[
 "position": 0, "word": "Le", "concept": null ,
 "position": 1, "word": "président", "concept": "president",
 "position": 2, "word": "préfère", "concept": "prefer-01",
 "position": 3, "word": "de", "concept": null,
 "position": 4, "word": "ne", "concept": ":polarity",
 "position": 5, "word": "pas", "concept": ":polarity",
 "position": 6, "word": "organiser", "concept": "hold-04",
 "position": 7, "word": "des", "concept": null,
 "position": 8, "word": "éléctions", "concept": "elect-01",
 "position": 9, "word": "anticipées", "concept": "early",
 "position": 10, "word": ".", "concept": null
], [
...
```

Figure 7: Prompt for GPT-4.1-mini to align tokens of a sentence with the AMR graph

An initial comparison produced expected results, for instance that a `nsubj` dependency relation (with verbs in Active Voice) corresponds mainly to an `:ARG0` relation in the AMR graph across the four languages which we examined. We counted the co-occurences of a dependency relation between two tokens and the AMR relation between instances aligned to the tokens (Table 1). For instance "this" is `nsubj` of "[be] different" in the (English) dependency tree in Fig. 6 (orange back-

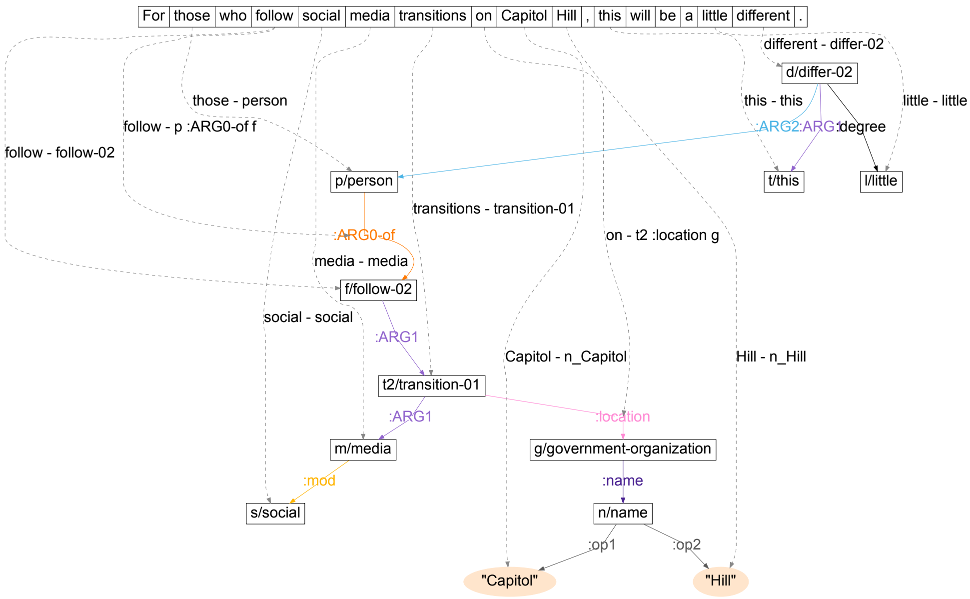

Figure 4: AMR graph and corresponding English tokens aligned (id: n01001013)

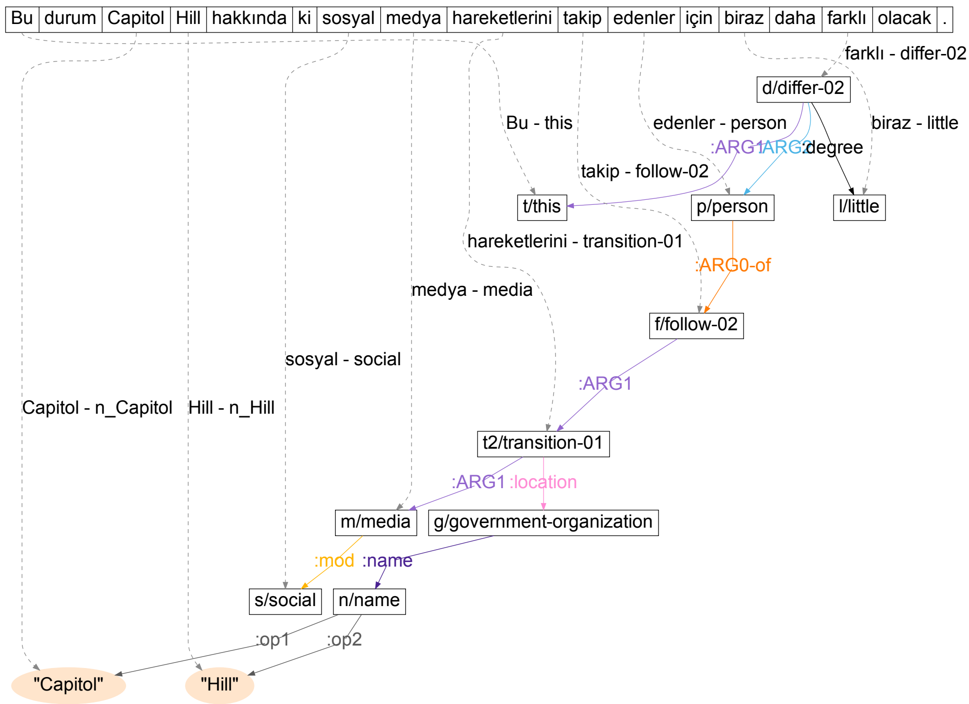

Figure 5: AMR graph and corresponding Turkish tokens aligned (id: n01001013). *Durum* ("situation") is not aligned since this is implicit in the English sentence which was used to create the AMR graph.

ground) and the tokens "this" and "`different`" of the syntax tree are aligned to the instances `t/this` and `d/differ-02` of the AMR graph, respectively (Fig. 4). So the dependency relation `nsubj` can be mapped to the AMR relation `:ARG0`. 36.1% of `nsubj` (279 examples in ParallelUD-AMR) can be mapped to `:ARG0` in English (cf. italic line in Table 1). The rates for French, Turkish and Arabic are lower but still the most frequent mappings for `nsubj`.

6. Conclusion and perspectives

We presented a new dataset ParallelUD-AMR, an Abstract Meaning Annotation of the the English corpus of PUD. We provided an initial mapping of dependency relations with AMR relations (based on tokens and their alignments within the AMR

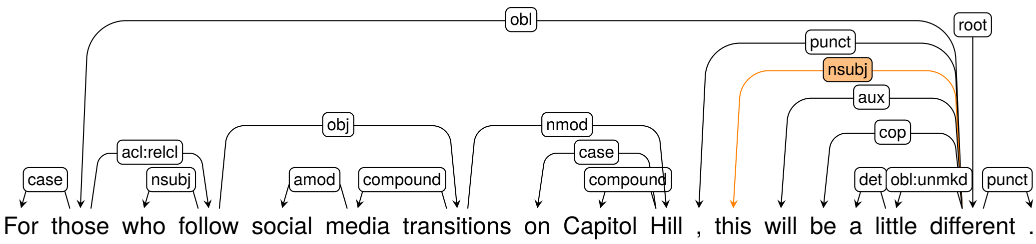

Figure 6: Dependency annotation of sentence n01001013 (UD_English-PUD)

deprel	AMR	English		French		Arabic		Turkish	
acl	:ARG1-of	33.3%	33	35.1%	26	†		19.0%	47
acl:relcl	:ARG0-of	29.3%	27	20.2%	21	16.0%	25	*	
acl:relcl	:ARG1-of	25.0%	23	9.6%	10	14.7%	23	*	
advcl	:time	12.1%	19	15.0%	16	7.3%	12	5.3%	12
advcl	:purpose	9.6%	15	6.5%	7	3.0%	5	2.7%	6
advmod	:mod	16.7%	76	7.9%	39	12.0%	28	7.1%	24
amod	:mod	31.0%	201	26.0%	174	19.1%	154	18.7%	127
ccomp	:ARG1	29.3%	53	36.1%	35	50.0%	43	41.6%	37
nmod	:ARG1	18.3%	88	11.9%	112	12.3%	191	5.3%	5
nsubj	*:ARG0*	*36.3%*	*279*	*29.1%*	*218*	*20.6%*	*145*	*23.2%*	*143*
nsubj	:ARG1	15.9%	122	9.6%	72	13.4%	94	18.3%	113
nsubj:pass	:ARG1	56.6%	43	44.0%	33	39.3%	22	*	
obj	:ARG1	67.1%	286	40.6%	217	30.3%	115	45.4%	210
obl	:ARG2	19.1%	104	11.3%	74	7.7%	74	7.5%	52
obl	:ARG1	8.3%	45	7.2%	47	10.9%	104	3.3%	23
xcomp	:ARG1	56.2%	81	34.4%	75	18.2%	14	28.4%	25
xcomp	:ARG2	16.0%	23	11.5%	25	4.5%	4	10.4%	8

Table 1: The most frequent corresponding dependency and AMR relations (list non exhaustive). Notes: †: few `acl` relations in Arabic-PUD; *: no `acl:relcl` nor `nsubj:pass` in Turkish-PUD. Note that alignments in Turkish and Arabic have not been manually validated.

graph). The mapping can certainly be further refined. For instance, we need to check whether `nsubj` - `:ARG1` mappings are simply because the AMR concept (PropBank) does not have an `:ARG0` but only an `:ARG1` or if the verb is in Mediopassive voice (cf. "the window opens"). Eventually the Feature-column of the UD data can be taken itno account.

Further interesting results could yield the mapping of an Ergative language like Basque (language isolate) or mixed Ergative-Accusative Georgian (Kartvelian language), or a Topic language like Tagalog or Bikol (both Austronesion languages). However the PUD corpus has currently not been translated into any of these languages.

Other possible usages of ParallelUD-AMR and PUD data are trying to find inconsistencies by identifying unexpected mappings like `amod` - `:ARG2` which might identify an invalid dependency or AMR-relation.

Finally other usages of ParallelUD-AMR together with the PUD treebanks are conceivable: Could ParallelUD-AMR assist Dependency parsers to provide better results? Or vice-versa, could AMR parsers exploit PUD treebanks to enhance their score? The syntactic-semantic annotation could also help to bootstrap AMR annotations from existing UD data.

7. Acknowledgements

We are grateful to the colleagues who participated in the annotation effort.

8. Bibliographical References

Alan Akbik, Laura Chiticariu, Marina Danilevsky, Yunyao Li, Shivakumar Vaithyanathan, and Huaiyu Zhu. 2015. Generating high quality proposition Banks for multilingual semantic role labeling. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Vol-*

ume 1: Long Papers), pages 397–407, Beijing, China. Association for Computational Linguistics.

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

Shu Cai and Kevin Knight. 2013. Smatch: an Evaluation Metric for Semantic Feature Structures. In *51st Annual Meeting of the Association for Computational Linguistics*, page 748–752, Sofia, Bulgaria. Association for Computational Linguistics.

Andrew Drozdov, Jiawei Zhou, Radu Florian, Andrew McCallum, Tahira Naseem, Yoon Kim, and Ramón Astudillo. 2022. Inducing and using alignments for transition-based AMR parsing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1086–1098, Seattle, United States. Association for Computational Linguistics.

Federica Gamba, Alexis Palmer, and Daniel Zeman. 2025. Bootstrapping UMRs from Universal Dependencies for Scalable Multilingual Annotation. In *Proceedings of the 19th Linguistic Annotation Workshop (LAW-XIX-2025)*, pages 126–136, Vienna, Austria. Association for Computational Linguistics.

Kelvin Han and Siyana Pavlova. 2019. Going from UD towards AMR. Technical report, Laboratoire Lorrain de Recherche en Informatique et ses Applications — UMR 7503.

Johannes Heinecke. 2023. metAMoRphosED: a graphical editor for Abstract Meaning Representation. In *19th Joint ACL - ISO Workshop on Interoperable Semantic Annotation at IWCS 2023*, Nancy. Tilburg University.

Johannes Heinecke and Anastasia Shimorina. 2022. Multilingual Abstract Meaning Representation for Celtic Languages. In *Proceedings of the 4th Celtic Language Technology Workshop within LREC2022*, pages 1–6, Marseille. ELRA.

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*.

Paul Kingsbury and Martha Palmer. 2002. From TreeBank to PropBank. In *Proceedings of the Third International Conference on Language Resources and Evaluation*, pages 1989–1993, Las Palmas, Canary Islands - Spain. European Language Resources Association.

Young-Suk Lee, Ramón Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian, and Salim Roukos. 2022. Maximum Bayes Smatch ensemble distillation for AMR parsing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5379–5392, Seattle, United States. Association for Computational Linguistics.

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

Simon Mille, Anja Belz, Bernd Bohnet, Yvette Graham, Emily Pitler, and Leo Wanner. 2018. The first multilingual surface realisation shared task (SR'18): Overview and evaluation results. In *Proceedings of the First Workshop on Multilingual Surface Realisation*, pages 1–12, Melbourne, Australia. Association for Computational Linguistics.

Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher Manning, Sampo Pyysalo, Sebastian Schuster, Tyers Francis, and Daniel Zeman. 2020. Universal Dependencies v2: An Evergrowing Multilingual Treebank Collection. In *LREC 2020*, Marseille. ELRA.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106.

Michael Pust, Ulf Hermjakob, Kevin Knight, Daniel Marcu, and Jonathan May. 2015. Parsing English into Abstract Meaning Representation Using Syntax-Based Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1143–1154, Lisbon, Protugal. Association for Computational Linguistics.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A Python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.

Sebastian Schuster and Christopher D. Manning. 2016. Enhanced English Universal Dependencies: An Improved Representation for Natural Language Understanding Tasks. In *the tenth international conference on Language Resources and Evaluation*, pages 2371–2378, Portorož, Slovenia. European Language Resources Association.

Milan Straka. 2018. UDPipe 2.0 Prototype at CoNLL 2018 UD Shared Task. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 197–207, Brussels. Association for Computational Linguistics.

Milan Straka and Jana Straková. 2017. Tokenizing, POS Tagging, Lemmatizing and Parsing UD 2.0 with UDPipe. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 88–99, Vancouver, Canada. Association for Computational Linguistics.

Ida Szubert, Adam Lopez, and Nathan Schneider. 2018. A structured syntax-semantics interface for English-AMR alignment. 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 1169–1180, New Orleans, Louisiana. Association for Computational Linguistics.

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 James Pustejovsky, Rosa Vallejos, and Nianwen Xue. 2021. Designing a Uniform Meaning Representation for Natural Language Processing. *Künstliche Intelligenz*, 35:343–360.

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.

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 and Nathan Schneider. 2024. Assessing the Cross-linguistic Utility of Abstract Meaning Representation. *Computational Linguistics*, 50(2):419–473.

9. Language Resource References

Damonte, Marco and Cohen, Shay. 2020. *AAbstract Meaning Representation 2.0 - Four Translations*. Linguistic Data Consortium. distributed via LDC: LDC2020T07, ISLRN 359-968-732-813-3. PID https://catalog.ldc.upenn.edu/LDC2020T07.

Kevin Knight and Bianca Badarau and Laura Baranescu and Claire Bonial and Madalina Bardocz and Kira Griffitt and Ulf Hermjakob and Daniel Marcu and Martha Palmer and Tim O'Gorman and Nathan Schneider. 2020. *Abstract Meaning Representation (AMR) Annotation Release 3.0*. Linguistic Data Consortium. distributed via LDC: LDC2020T02, 3.0, ISLRN 676-697-177-821-8. PID https://catalog.ldc.upenn.edu/LDC2020T02.

Zeman, Daniel et al. 2025. *Universal Dependencies 2.16*. PID http://hdl.handle.net/11234/1-5901. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL).