

Modeling Quantification and Scope in Abstract Meaning Representations

James Pustejovsky, Nianwen Xue, Kenneth Lai

Department of Computer Science
Brandeis University
Waltham, MA USA

{jamesp,xuen,klai12}@brandeis.edu

Abstract

In this paper, we propose an extension to Abstract Meaning Representations (AMRs) to encode scope information of quantifiers and negation, in a way that overcomes the semantic gaps of the schema while maintaining its cognitive simplicity. Specifically, we address three phenomena not previously part of the AMR specification: quantification, negation (generally), and modality. The resulting representation, which we call “Uniform Meaning Representation” (UMR), adopts the predicative core of AMR and embeds it under a “scope” graph when appropriate. UMR representations differ from other treatments of quantification and modal scope phenomena in two ways: (a) they are more transparent; and (b) they specify default scope when possible.

1 Abstract Meaning Representations

Abstract Meaning Representations (AMRs) have recently become popular as a strategy for encoding a kind of canonical meaning for natural language sentences (Banarescu et al., 2013). They differ significantly from other encoding schemes used in NLP—e.g., minimal recursion semantics (MRS)—in terms of their expressiveness for several semantic phenomena in natural language (Copestake et al., 2005). Still, in spite of such shortcomings, there is a major attraction to the general philosophy of this approach: by focusing on the predicative core of a sentence, it is an intuitive representation for both interpreting the semantics of a sentence, and perhaps more importantly, for use in annotation efforts.


An AMR represents the meaning of a sentence with a single-rooted, directed, acyclic graph with nodes labeled with concepts and edges labeled with relations. The primary component of an AMR is the predicate-argument structure, with the predicate being a concept that takes a number of

arguments as its children. The predicate and its arguments are represented as nodes in the AMR graph, and the edges represent the relation between the predicate and each of its arguments. As an illustration, the PENMAN notation and graph representation below in (2) represent the AMR for the sentence in (1).

(1) John can’t afford a car at the moment.

(2) a. (p / possible-01
:ARG0 (a / afford-01
:ARG0 (p2 / person
:name (n / name
:op "John"))
:ARG1 (c / car)
:time (m / moment))
:polarity -)

b.



Propositions in an AMR are sense-disambiguated (Palmer et al., 2005). In the example above, “possible-01” refers to the first sense of “possible” while “afford-01” represents the first sense of “afford”. A predicate can take a number of core arguments (ARG0, ARG1, etc.) as well as adjunct arguments (e.g., *time*). The semantic roles for the core arguments are defined with respect to each sense of a predicate

and they are drawn from the PropBank frame files ¹. For example, the semantic roles for the core arguments of different senses of “afford” are defined as follows:

- (3) a. **afford-01**: be able to spare, have the financial means
 ARG0: haver of financial means, agent
 ARG2: costly thing, theme
- b. **afford-02**: provide, make available
 ARG0: provider, agent
 ARG1: provided, theme
 ARG2: recipient

The attraction of AMR-style representations and annotations is the adoption of a *predicative core* element along with its arguments: e.g., an event and its participants. This, in turn, leads to an event-rooted graph that has many advantages for parsing and matching algorithms. As can be seen from the example, the predicate-argument structure is front and center in AMR, and we consider this to be one of its strengths.

However, as it currently stands, AMR does not represent quantification or its interaction with modality and negation (Bos, 2016). The challenge is to maintain the focus on the predicate-argument structure while also adequately accounting for linguistic phenomena that operate above the level of the core AMR representation, in particular quantification and modality.

2 Quantification and Scope

It can be argued that, besides graph-based matching over predicative structures, AMR does not provide good support for logical inference because it does not yet properly handle scoping and other phenomena. For example, in (4), there is a single talk that everyone in the room is listening to, while in (5), each person has their own coffee. However, AMR does not distinguish between these two cases: it could just as well be that everyone in the room listened to a different talk, or that everyone at noon shared a single cup of coffee.

- (4) a. Everyone in the room listened to a talk.
 b. $\exists y[\text{talk}(y) \wedge \forall x \exists e[\text{person}(x) \wedge \text{inRoom}(x) \rightarrow \text{listen}(e, x, y)]]$
 c. (l / listen-01
 :ARG0 (p / person
 :mod (a / all)
 :location (r / room))
 :ARG1 (t / talk))

- (5) a. Everyone drank a coffee at noon.
 b. $\forall x[\text{person}(x) \rightarrow \exists y \exists e[\text{coffee}(y) \wedge \text{drink}(e, x, y) \wedge @ (e, \text{noon})]]$
 c. (d / drink-01
 :ARG0 (p / person
 :mod (a / all))
 :ARG1 (c / coffee)
 :time (n / noon))

In fact, this inability of AMRs to distinguish scoping relations among quantifiers also extends to negation and modality. For example, the AMR for the sentence “Every student did not fail” is given below.

- (6) a. (d / fail-01
 :ARG0 (s / student
 :mod (a / all))
 :polarity (-))
 b.
-

The sentence is ambiguous, however, between the readings “for every student, that student did not fail” and “it is not the case that every student failed”.

While MRS and other flattened semantic representations provide a solution to these issues, giving faithful translations of scope with typed expressions, there are several drawbacks to these approaches. Flat representations reveal no semantic core. Hence, as annotations, the resulting structures are difficult to interpret and inspect. Furthermore, quantifier scope is often underspecified even when it can be disambiguated in context. Dependency MRS (DMRS) is one exception to this in the MRS family of representations (Copestake, 2009), where dependency relations link argument heads to the major predicator of the sentence.

In our research, we propose to represent scope relationally, while maintaining both the centrality of the predicative core of the sentence (e.g., *listen*, *drink*), as well as the syntactic integrity of the quantified expression (e.g., *every person*). A relational interpretation for scope provides a first-order interpretation: it references two specific nodes in the graph, and orders one relative to the other. This operates over generalized quantifiers (*some book*, *most people*), negation (*not*, *no*), as well as modals (*possibly*, *likely*, *must*). From an

¹<https://verbs.colorado.edu/verb-index>

annotation perspective, this is quite different from flat structures, since a human judgment in scope between two elements is directly reflected in the resulting graph. There are complex interactions between negation, modal expressions, and quantified NPs that we will examine, first representationally, and then experimentally with small-scale annotation and testing.

We believe there are advantages to adopting an AMR-style representation for predicate-argument forms of sentences (Banarescu et al., 2013). Given the complexity inherent in the semantics of number, negation, and quantification, we believe that a similar approach to the annotation of scope has some advantages. These include the following:

- It maintains a focus on the *predicative core* of the sentence;
- There is likely a lower *cognitive load* for annotation by non-experts;
- Semantic relations are *transparent* in the graphical representation.

Addressing the problems associated with scope adopting this approach results in a representation we call “Uniform Meaning Representation” (UMR), where the predicative core of AMR is maintained, and embedded under a “scope” graph when required.

3 Towards a Uniform Meaning Representation for Scope

In this section, we illustrate our approach to encoding the expression of quantifier scope in UMR. We draw on some work within the ISO annotation community, where the problem of explicitly annotating scoping relations of events and temporal or spatial quantifiers has been addressed.

To explicitly represent relative scope of quantified expressions, ISO-Space (Pustejovsky, 2017) uses the @quant attribute (adopted from ISO-TimeML), applying it to spatial entities, and in addition uses the attribute @scopes to specify a scoping relation. The following example, taken from ISO 24617-7:2014, illustrates this:

- (7) a. A computer_{se1} is on_{ss1} every desk_{se2}.
 b. <spatialEntity id="se1" pred="computer"
 quant="1" scopes="" />
 <spatialEntity id="se2" pred="desk"
 quant="every" scopes="#se1"/>

From a semantic point of view, however, this use of the @scopes attribute is unsatisfactory since the relative scoping of quantifications over different sets of entities is not a local property of one of these quantifications; therefore an annotation such as (7) does not have a compositional semantics. Therefore, we follow (Bunt et al., 2018) and use a link structure, **scopeLink**, to represent scope relations among quantifying NPs, where relType takes a value of ‘narrower’, ‘wider’, or ‘equal’. For the example in (7), this amounts to marking the universal as taking wide scope over the indefinite.

- (8) a. scopeLink(*arg*₁, *arg*₂, *relType*)
 b. <scopeLink arg1="#se2" arg2="#se1"
 relType="wider"/>

We modify this scoping relation by introducing the predicative domain as an additional argument,

- (9) a. $\lambda pred \lambda a_1 \lambda a_0 [\mathbf{scope}(a_0, a_1, pred)]$

and model the semantic effect of this relation as similar to the mechanism of Cooper Storage (Cooper, 1975) or a continuation-passing style interpretation of generalized quantifiers (Barker, 2002).

$\llbracket a_0 \rrbracket (\llbracket a_1 \rrbracket (\llbracket pred \rrbracket))$

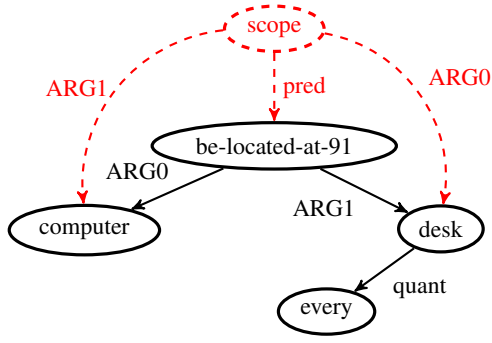
For example, consider a relation with two quantifier phrases, and a scoping of the direct object QP_2 over QP_1 :

- (10) a. $[QP_1]_{arg0} \text{ pred } [QP_2]_{arg1}$
 b. **scope**(QP_1, QP_2, pred)

The ordering of arguments determines the function application order of each expression, as with continuation-passing style.

This representation is convenient, in that we can maintain a rooted graph structure with the **scope** relation as the root node, as demonstrated below.

- (11) a. (s / scope
 :pred (b / be-located-at-91
 :ARG0 (c / computer)
 :ARG1 (d / desk
 :quant (e / every)))
 :ARG0 d
 :ARG1 c)
 b.



Given the scope-rooted graph above, we apply an interpretation function that translates the graph to an appropriate first-order expression. This gives:

$$(12) \llbracket \text{every desk} \rrbracket (\llbracket \text{a computer} \rrbracket (\llbracket \text{be_located} \rrbracket))$$

For the present annotation, we arrive at the expression in (13).

$$(13) \forall y[\text{desk}(y) \rightarrow \exists x[\text{computer}(x) \wedge \text{be-located-at}(x, y)]]$$

With the introduction of scope over quantifiers, the annotation provided by a UMR can be compared more directly to the approach and representations deployed in the Groningen Meaning Bank (GMB) and the Parallel Meaning Bank (PMB) projects (Bos et al., 2017; Abzianidze et al., 2017; Van Noord et al., 2018). In this work, sentences are expressed as DRSs within Discourse Representation Theory (Kamp and Reyle, 1993). However, most of the sentences in PMB with potential quantifier scope ambiguities involve temporal expressions and their relative scope over event variables, rather than quantified arguments to the verb. An example is that shown in (14).

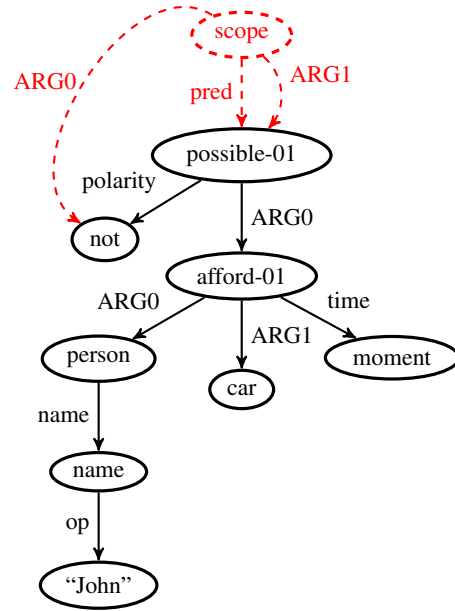
$$(14) \text{ a. John golfed every Sunday.} \\ \text{ b. } \forall t[\text{Sunday}(t) \rightarrow \exists e[\text{golf}(e, j) \wedge \text{on}(e, t)]]$$

The strategy taken by (Bos et al., 2017), followed here as well, is to scope temporal expressions over the events they govern.

Now let us see how the **scope** relation can be deployed to handle negation and modality in UMR. Consider first the treatment of modals in AMR. As seen in (2) above, modals are treated as predicative nodes. Hence, from $(p / \text{possible-01} : \text{ARG0 } \phi)$, we can derive the equivalent propositional modal expression, $\diamond\phi$. However, in (2) we need to translate the polarity over the modal appropriately: $\neg\diamond\phi$.

In UMR, the **scope** relation acts as a root node assigning the polarity value as taking scope over the modal, along with its body. Consider the UMR graph as shown below. Note that because there may be multiple negations in a sentence, we index negations, e.g., $(n2 / \text{not})$.

$$(15) \text{ a. } (s / \text{scope} \\ : \text{pred } (p / \text{possible-01} \\ : \text{ARG0 } (a / \text{afford-01} \\ : \text{ARG0 } (p2 / \text{person} \\ : \text{name } (n / \text{name} \\ : \text{op } \text{"John"})) \\ : \text{ARG1 } (c / \text{car}) \\ : \text{time } (m / \text{moment})) \\ : \text{polarity } (n2 / \text{not})) \\ : \text{ARG0 } n2 \\ : \text{ARG1 } p) \text{ b.}$$



The graph-interpretation function continues walking down the tree, and expands the Skolemized form for ‘car’ into a quantified expression, inside the scope of the modal, as shown below.

$$(16) \neg\diamond[\exists x[\text{car}(x) \wedge \exists e[\text{afford}(e, j, x) \wedge @\langle e, N \rangle]]]$$


This can be compared to the first-order modal expression generated by (Bos, 2015; Bos et al., 2017) for the sentence as shown below in (17).

$$(17) \neg\exists x[\text{car}(x) \wedge \diamond\exists e[\text{afford}(e, j, x) \wedge @\langle e, N \rangle]]$$

Thus far we have briefly examined the following semantic constructions: quantifier scope for arguments; temporal adjuncts over events; and relative scope of negation and modality.

Now consider the interaction of negation with quantifiers in AMR, as seen in the possible interpretations of (18).

(18)



There are two interpretations consistent with this AMR graph.

(19) a. There is no event of a passenger surviving (no one survived).


$$\neg\exists e\exists x[\textit{survive}(e, x) \wedge \textit{person}(x)]$$

b. There is a passenger who did not survive.

$$\exists x[\textit{person}(x) \wedge \neg\exists e[\textit{survive}(e, x)]]$$

With the introduction of the **scope** relation node, we can distinguish these interpretations: for example, the reading in (19a) would be represented as shown below.

(20)



4 Conclusion

In this short note, we introduced a representation and interpretive strategy for capturing scope relations between quantifiers, negation, and modals in AMR. This required an enrichment to the basic vocabulary of AMR that we refer to as a *Uniform Meaning Representation*. The UMR strategy adopts one of the more attractive features of AMR, the predicative core, while increasing the representation language’s expressive coverage with the introduction of a **scope** node, determining the relative scope between its two arguments. The interpretation of a specific quantifier or modal is lexically determined. This work is part of a combined effort to enrich the representation of AMRs with tense, (Donatelli et al., 2019), discourse relations (O’Gorman et al., 2018), quantification, and modality.

Acknowledgement

We would like to thank the anonymous reviewers for their helpful comments. This work is

supported by the IIS Division of National Science Foundation via Award No. 1763926 entitled “Building a Uniform Meaning Representation for Natural Language Processing”. All views expressed in this paper are those of the authors and do not necessarily represent the view of the National Science Foundation.

References

- 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. *arXiv preprint arXiv:1702.03964*.
- 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.
- Chris Barker. 2002. Continuations and the nature of quantification. *Natural language semantics*, 10(3):211–242.
- Johan Bos. 2015. Open-domain semantic parsing with boxer. In *Proceedings of the 20th nordic conference of computational linguistics (NODALIDA 2015)*, pages 301–304.
- Johan Bos. 2016. Expressive power of abstract meaning representations. *Computational Linguistics*, 42(3):527–535.
- Johan Bos, Valerio Basile, Kilian Evang, Noortje J Venhuizen, and Johannes Bjerva. 2017. The groningen meaning bank. In Nancy Ide and James Pustejovsky, editors, *Handbook of linguistic annotation*, pages 463–496. Springer.
- Harry Bunt, James Pustejovsky, and Kiyong Lee. 2018. Towards an iso standard for the annotation of quantification. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- Robin Hayes Cooper. 1975. Montague’s semantic theory and transformational syntax.
- Ann Copestake. 2009. Slacker semantics: why superficiality, dependency and avoidance of commitment can be the right way to go. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 1–9.
- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A Sag. 2005. Minimal recursion semantics: An introduction. *Research on language and computation*, 3(2-3):281–332.
- Lucia Donatelli, Nathan Schneider, William Croft, and Michael Regan. 2019. Tense and aspect semantics for sentential amr. *Proceedings of the Society for Computation in Linguistics*, 2(1):346–348.
- Hans Kamp and Uwe Reyle. 1993. From discourse to logic, volume 42 of studies in linguistics and philosophy.
- 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.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.
- James Pustejovsky. 2017. ISO-Space: Annotating static and dynamic spatial information. In *Handbook of Linguistic Annotation*, pages 989–1024. Springer.
- Rik Van Noord, Lasha Abzianidze, Hessel Haagsma, and Johan Bos. 2018. Evaluating scoped meaning representations. *arXiv preprint arXiv:1802.08599*.