Quantification in Abstract Meaning Representation

Kiyong Lee and Chongwon Park
Korea University, Seoul, University of Minnesota Duluth
ikiyong@gmail.com, cpark2@d.umn.edu

Abstract

Quantification is common in text but underrepresented in AMR. It also strains the common conjunctive interpretation of AMR graphs, since universal quantification introduces scope-taking structure and variable binding that cannot be captured as a flat list of conjuncts. We propose an enriched AMR that supports quantificational meaning while preserving AMR's graph backbone. At the predicate level, we add QuantML-style features such as domain restriction, determinacy, distributivity, and involvement. At the discourse level, we add contextual constraints that encode scope and other discourse-sensitive conditions. The two levels follow the UMR architecture and are linked by shared identifiers. We map the enriched graphs to two-block logical forms: a minimal predicate structure of events and participants, plus a constraint block that relates them.

Keywords: QuantML, quantification, minimal logical form, predicate structure, discourse-level contextual constraint

1. Motivation, Aim, and Tasks

Since Montague (1974) and Bach et al. (1995), quantification has been a central and contested topic in natural language semantics. Quantificational expressions are pervasive in ordinary language and raise some of the most technically demanding problems in semantic analysis, including scope ambiguity, distributivity, and domain restriction. By contrast, AMR work has largely set aside quantification, a gap that can be detrimental, especially for downstream applications that require precise meaning representations. A further motivation is that, under a common baseline interpretation where an AMR corresponds to a conjunction of atomic conditions, universal quantification cannot be handled as just another conjunct, because it requires a conditional (implication) structure and explicit variable binding, giving rise to well-known scope and bound-variable problems (Bos, 2016, 2020).

This paper investigates how quantificational phenomena, including scope ambiguity and discourse-sensitive constraints on the interpretation of quantified expressions, can be represented by enriching Abstract Meaning Representation (AMR) with a quantification layer compatible with Uniform Meaning Representation (UMR). We adopt the ISO 24617-12:2025 QuantML feature inventory, including domain restriction, determinacy, distributivity, and involvement, and encode these features as annotations that can be mapped into enriched AMR structures.

For this purpose, this paper undertakes three specific tasks. The first task is to adopt the two-layer architecture of Uniform Meaning Representation (UMR, 2022), pairing a sentence-level AMR-style predicate-argument graph with a document-level component that encodes cross-sentential relations such as temporal and modal dependencies and coreference. The document-level component is linked to sentence-level nodes through shared identifiers (for example, by referencing sentence-indexed concept ID's), yielding an integrated representation that can support discourse-sensitive constraints on quantifier interpretation.

The second task is to define a set of discourse-sensitive constraints in the document-level component. The set includes constraints inferred from context and those triggered by quantificational features annotated at the sentence-level AMR nodes. The constraint formalism is inspired by Bos (2020)'s AMR$^+$ proposal, which preserves AMR's predicate argument backbone. We do not, however, follow Bos's way of marking those constraints by adding a logical layer with explicit indices.

The third task is to construct two-block logical forms, roughly corresponding to the two-level serialization of each annotation structure. The first block represents a (minimal) predicate structure, consisting of an event and a list of objects that participate in it, while the second block represents how all of the logical forms in the first block are related by contextual discourse constraints.

2. Main Issues

Two clarifications are needed before turning to the technical issues in this section. The first concerns why we keep the primary meaning representation graph-based and comparatively simple. The second concerns why we add an explicit logical interpretation layer rather than reading quantificational scope directly from the graph.

First, we treat enriched AMR as the primary interface because the paper is aimed at a represen-

Proceedings of the Workshop on Structured Linguistic Data and Evaluation, pages 104–113
May 11, 2026.

tation suitable for corpus annotation and computational pipelines, not only for hand-built semantic analyses. AMR's practical value has been its compact predicate-argument graph backbone and its compatibility with existing annotation practice and parsers (Banarescu et al., 2019; Knight et al., 2020). If quantificational information is introduced in a way that forces pervasive reification, explicit variable bookkeeping, or global scopal decisions at every step, the representation becomes harder to annotate consistently and harder to learn from data. For this reason, we aim for a conservative extension strategy. That is, quantificational content is added as a small inventory of local features and discourse-level constraints that remain compatible with the two-layer UMR architecture (UMR, 2022) and that can be mapped from the QuantML annotation scheme and inventory (Bunt et al., 2018, 2022; ISO, 2025). This design keeps the graph readable, keeps the enrichment optional for applications that do not need it, and preserves interoperability with existing AMR tooling.

Second, the logical forms in this paper are neither a competing representation layer nor an additional annotation target. They are used as a reference semantics for the enriched graphs. Once the representation is enriched with meaning-bearing features and discourse constraints, readers need an explicit statement of what those annotations commit us to in interpretation. The logical forms provide that statement in a standard, checkable metalanguage. They make scope, domain restriction, and distributive conditions explicit enough to (i) verify that the enriched graphs distinguish the intended readings, (ii) validate conversions between QuantML-style annotation structures and AMR-style graphs, and (iii) support accuracy claims about the representation by comparing whether two analyses yield the same or different truth-conditional commitments (Bos, 2016, 2020).

A key reason for spelling out the reference semantics is that a common baseline interpretation of AMR treats an AMR graph as a flat conjunction of atomic conditions (a list of conjuncts). This works well for many positive, existential-like cases, but it breaks down once scopal operators are needed. In particular, universal quantification requires a conditional (implication) structure, which introduces explicit variable binding and resolves the bound-variable and scope problems that arise under a purely conjunctive view (Bos, 2016, 2020). An equivalent option is to compile universals into negation and existential quantification (using $\forall x\, \phi \equiv \neg \exists x\, \neg \phi$), but this still requires scope-taking operators and therefore cannot be recovered from conjunction alone (Bos, 2016, Section 5). The logical forms in this paper make these scope and binding commitments explicit while keeping the annotation-layer graphs lightweight.

2.1. Simplicity vs. Adequacy

The first design issue concerns the balance between notational simplicity and representational adequacy in meaning representation. This section compares two concrete syntaxes for encoding quantificational analyses: the XML-based representation format, used in ISO QuantML (ISO, 2025), and the PENMAN notation, used to serialize AMR-style rooted, directed graphs.

Within the ISO 24617 Semantic Annotation Framework, QuantML specifies an XML-based representation format for annotation structures. This design choice follows the broader ISO practice of using XML as a uniform concrete syntax across annotation standards. It is largely for syntactic interoperability and reuse of existing XML tooling, rather than for compactness or ease of direct human authoring. At the same time, QuantML distinguishes an abstract syntax from its XML concrete syntax, which in principle allows alternative serializations that preserve the same underlying annotation structures (Bunt, 2010; ISO, 2016).

For dense semantic annotation and integration with AMR-style graphs, however, XML can be prohibitively verbose. Even for small fragments, an XML encoding like Example (1) requires nested elements, semantic categories encoded as tag names, and multiple identifiers for cross-references. These devices make entity structures and link structures explicit, but they also obscure the underlying graph and complicate both manual inspection and conversion into a graph-based meaning representation.

Consider Example (1) that annotates a fragment of text in XML:

(1) Sentence: Dogs bark.
Segmented: Dogs$_{w1}$ bark$_{w2}$.
```
<ENTITY xml:id="x1"
target="#w1" pred="dog"/>
<EVENT xml:id="e1" target="#w2"
pred="bark"/>
<SRLINK eventID="#e1"
participantID="#x1"
relType="agent"/>
```

In contrast, AMR encodes the same semantic content for the purposes of sentence interpretation, but does so in a more compact and direct notation.

(2) :: Sentence: Dogs bark.
:: Interpretation: For the present comparison, the example introduces a barking event with a dog participant.
:: id: amr001KL,
:: Data-creation time: 2026-01-23
```
(b / bark-01
   :agent (d / dog))
```

Although the surface format differs substantially, the AMR representation in Example (2) introduces the same core semantic objects and relations: an eventuality of barking and a dog participant linked to that event. As a result, both representations support the same core predicative interpretation of the sentence, despite their different levels of structural and notational complexity.

By comparing the two representations, (1) and (2), we observe that several components of the XML representation can be removed to reduce the size of the AMR representation. First, all of the tag names, such as `ENTITY, EVENT`, and `SRLINK`, need not be represented in AMR. Second, the link for participant-role assignment is unnecessary; it is built into the AMR structure.

2.2. Logical Forms and their Interpretation Model

The second issue concerns how to interpret enriched AMR structures for quantification.

(3) :: Sentence: All dogs bark.
:: Interpretation:
There are dogs, and all of them bark.
:: id: amr002KL,
:: Data-creation time: 2026-03-22

```
(b / bark-01
   :agent (d / dog
             :num PL
             :quant A))
```

Here, PL and A are logical constants standing for plurality and universal quantification. A minimal predicate structure is obtained by abstracting away from these two feature specifications.

2.2.1. Minimal Logical Form Representing a Predicate Structure

To this end, we first derive logical forms directly from minimal AMR predicate structures, such as (4).

(4) Minimal Logical Form
$\{b, d\}$
$[instance(b, bark),$
$instance(d, dog),$
$agent(b, d)]$

In AMR, b and d are treated as variables, but they can be understood as discourse referents in DRS (Kamp and Reyle, 1993). The comman (,) stands for logical conjunction ($\wedge$). The relation $instance$ or i, used in some places instead, is a convenient metalanguage predicate that links a discourse referent to an event predicate such as *bark* or an entity predicate such as *dog*.[1]

This is a *minimal logical form*, representing a predicate structure with its participants. It is minimal in two senses: first, it abstracts away from the feature specifications involving PL for plurality and A for universal quantification; second, it supports the construction of a minimal interpretation model in the spirit of Bos (2003) and Blackbun and Bos (2003); Blackburn and Bos (2004), based on the predicate structure contributed by AMR.

2.2.2. Minimal Model for Interpretation

In set-theoretic terms, a minimal interpretation model based on the AMR predicate structure (3), abstracted away from `:num PL` and `:quant A`, can be constructed as a tuple $<D, v>$, such that

(5) D=$\{b, d\}$
$v(instance)$=$\{(b, bark), (d, dog)\}$
$v(agent)$=$\{(b, d)\}$

The discourse domain D has two discourse referents, b and d. The semantic values of the two relations $instance$ and $agent$ are each specified as a set of pairs. This is a minimal model that satisfies each formula in (4), and thus the entire conjunctive form. It supports only the described situation: a barking event and a dog participant, with no further commitments beyond what the graph introduces.

Intuitively, the resulting minimal model describes a very small world, which we call the *described situation* (Barwise, 1981). The model is minimal in the sense that it does not commit us to any additional irrelevant individuals or facts beyond those required by the AMR structure itself. This paper adopts such a model construction for semantic interpretation.

2.2.3. Extension

The minimal model presented in (5) may appear trivial, but it is not fixed in a closed world. It may be extended when later discourse, contextual updates, or additional constraints require new individuals or additional relations. For example, suppose a situation is described in which two people are each drinking a beer. If subsequent discourse introduces the possibility that more people arrive and more beer is ordered, the minimal model can be extended by adding more people and more beer. In this way, the model functions as a small, incrementally extensible reference structure for interpretation.

Consider a simple example:

---

[1] Here $instance$ or i is used only in the reference semantics. It links discourse referents to the event or entity concepts introduced in the graph.

(6) :: Sentence: Three dogs are barking.
:: Interpretation: There are three dogs, and they are barking.

```
(b / bark-01
   :agent (d / dog
             :num 3)
   :aspect prog)
```

We construct a minimal model form by abstracting away from `:aspect prog` while retaining the cardinality information contributed by `:num 3`.

(7) Logical Form $\{b, x\}$
$[instance(b, bark),$
$instance(x, dog),$
$agent(b, x),$
$num(x)$=3]

Here, $num(x)$=3 is a short-hand representation, as defined:

(8) For any variable x, natural number n, and non-emptyset X,
$num(x) \geq n =_{def} x \in X$ such that $|X| \geq n$.

To interpret Logical Form (7), a minimal model is constructed as a tuple $<D, v>$, where D is the discourse domain and v assigns a semantic value to each of the binary relations, $instance$ and $agent$, such that we have:

(9) Extended Minimal Model
$D = \{b\} \cup \{x|x \in X, |X|$=3},
$v(instance)$=$\{(b, bark), (x, dog)\}$,
$v(agent)$=$\{(b, x)\}$

The variable x ranges over the set X, the cardinality of which is 3. This model thus supports a situation in which there are three agents, which are dogs, that bark.

2.2.4. Two-Block Representation of Logical Forms for Contextual Constraints

The proposed framework, in which AMR is enriched with UMR-style contextual constraints, yields a two-block representation for logical forms in cases (at least) involving quantification. Each AMR-based predicate structure contributes a minimal logical form in Block 1. In contrast, the discourse-level constraints are stated as quantified logical forms in Block 2. To interpret the resulting representation, we treat Block 1 (minimal structure) as a reference basis: formulas in Block 2 may contain variables that are not explicitly bound there, but that are intended to refer back to discourse entities introduced in Block 1. Operationally, such variables are resolved via Block 1 and bound at the appropriate scope in Block 2. In this respect, the overall structure is close to DRS (Kamp and Reyle, 1993), which combines a set of atomic conditions with a quantificational component.

(10) Restricted Discourse Domain:
A minimal model M_q for quantification is a tuple $<D, r, v>$, where
D is a non-empty domain of discourse, r is a restriction on D such that the restricted domain $D_r \subseteq D$, and v is a semantic valuation.

Restricted quantificational domains are introduced by the relevant noun phrase or, more often, contextually. The valuation v is a denotation function that satisfies the concept instances and relations contributed by the predicate-focused AMR structure.

(11) Logical Form:
Block 1 (AMR-based minimal structure):
$\{e, x\}$
$[instance(e, protest),$
$instance(x, student),$
$num(x) \geq 2,$
$dom(x) \cap r(x),$
$agent(e, x)]$
Block 2 (contextually constrained quantificational structure):
$\forall y$[y=x→
$\exists e_1[e_1 \sqsubseteq e \wedge agent(e_1, y)]]$

The representation in (11) lists the discourse entities $\{e, x\}$ that populate the minimal model, together with the conditions contributed by the AMR structure. The minimal model itself does not bind these variables; it records them as available referents. Block 2 then states the quantificational constraint over those referents.

The identity y=x in Block 2 is a *reference identifier*: It links Block 2 to the occurrences of x in Block 1, Predicate-focused minimal logical form. As defined in (8), the variable x is a variable in a set X the cardinality of which is 3. Hence, the quantificational form in Block 2 in (11) be rephrased as in (12) below:

(12) $\exists e[i(e, protest),$
$\forall x[[i(x, student), x \in X, |X| \geq 2] \rightarrow$
$\exists e_1[e_1 \sqsubseteq e, agent(e_1, x)]]],$
where the relation i again stands for the relation $instance$.

The effect of resolving the references to Block 1 can be made explicit by rewriting Block 2 so that the relevant discourse referents are existentially introduced along with the conditions that characterize them in the minimal model. This yields a compiled form such as (13):

(13) Block 2: Quantificational Structure:
$\forall y[y$=x
$\exists\{e_1, e\}[e_1 \sqsubseteq e \wedge i(e, protest) \wedge agent(e_1, y)]]$

Either representation corresponds to the same interpretive idea. One may interpret the original two-block structure by allowing Block 2 to refer back to Block 1, or one may compile that reference step into a single quantified formula as in (13). The two-block format is useful because it separates (i) the AMR-contributed described situation from (ii) the additional quantificational constraints imposed by the discourse level.

The distributive effect can be made more explicit by quantifying over event parts and requiring each member of the student plurality to participate in an appropriate subevent:

(14) Logical Forms:
a. Block 1: Preamble
$\{e, x\}$
$[i(e, protest),$
$i(x, student),$
$num(x) \geq 2,$
$dom(x) \cap r(x),$
$agent(e, x)]$
b. Block 2, referring to Block 1:
$\forall y[y{=}x \rightarrow \forall e_1[e_1 \sqsubseteq e \rightarrow agent(e_1, y)]]$

The logical form $dom(x) \cap r(x)$ in Block 1 represents the restricted domain of x for the universal quantification in Block 2. The quantified logical form in (b) captures the distributive effect by quantifying over event parts e_1 such that $e_1 \sqsubseteq e$ and by requiring $agent(e_1, y)$ for each such part. As before, the variables x and e are supplied by the discourse entities introduced in the first block.

In Section 3, we further discuss how this conversion can be triggered.

3. Constructing the Architecture of AMR Linked to UMR

Uniform Meaning Representation (UMR (2022), 0.9 Specification) extends AMR by pairing a sentence-level graph with a document-level graph that supports cross-sentential phenomena. UMR retains AMR's predicate-argument structures at the sentence level, but augments them with document-level annotations that link events and entities across sentences and encode temporal, coreferential, and modal information. Lee et al. (2025) suggest calling this UMR *document level* a *discourse level*, since many of these relations are defined over broader context rather than within a single sentence and can interact with phenomena such as scope and temporal interpretation.

Building on Pustejovsky et al. (2019), which addressed quantification and scope in AMR, we propose to join two levels: a predicate-structure level (P-level) and a discourse level (D-level). These two levels are connected by a root node, here called *Meaning*, as shown in Figure 1. The resulting representation is a single-rooted, labeled, directed graph with reentrancies; like standard AMR, it need not be a tree, as shown in Figure 2.

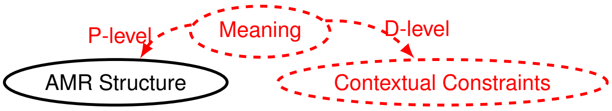

Figure 1: Architecture of the Enriched AMR

The syntax of both the predicate and discourse levels is stated in a PENMAN-style format compatible with current AMR conventions (Knight et al., 2020). We thus retain the same type of syntax for both levels. In the graph, each concept c in C, including the root, corresponds to a node. Relations label directed edges between nodes, where each edge goes from a parent node to a child node. A single parent node may have multiple children. The overall structure is interpreted as a labeled, directed graph whose connectivity is determined by these parent–child relations.

In the example below, each parenthesized variable introduces a node, and each colon-prefixed label (e.g., `:agent`, `:scope`) introduces a directed relation from the current node to its child, whether at the P-level or the D-level.

(15) :: Sentence: All the students protested.
:: Interpretation: There were students, and all of them participated in a protest.
:: id: 001, 2026-02-05

```
(r / meaning
   :P-level
     (e / protest-01
       :agent (x / student
                :num PL
                :dom (d / restr)
                :quant A))
   :D-level
      (c / constraints
         :domain (ex / entities
                    :op1 e
                    :op2 x)
         :scope (co / conditional
                    :op1 x
                    :op2 e)))
```

At the discourse level, constraints for interpretation specify the domain of discourse and the scopal relation of universal quantification. First, the domain lists the event variable e and the plural participant variable x. Second, the scope relation states that the student plurality takes scope with respect to the protest event. The corresponding two-block logical form can be stated as follows:

(16) a. $\{e, x\}$
$[instance(e, protest),$

$instance(x, student),$
$num(x) \geq 2,$
$dom(x) \cap r(x)]$
b. $\forall y[y{=}x \rightarrow$
$\exists e_1[e_1 \sqsubseteq e \wedge agent(e_1, y)]]$

Here, $num(x) \geq 2$ is shorthand for plurality: x denotes a plural discourse domain or for the range of x with at least two members.

The detailed annotation at the predicate and discourse levels can be incorporated into the general graph structure of the enriched AMR linked to the discourse level of UMR, following Pustejovsky et al. (2019) that combined two levels to treat scope ambiguity.

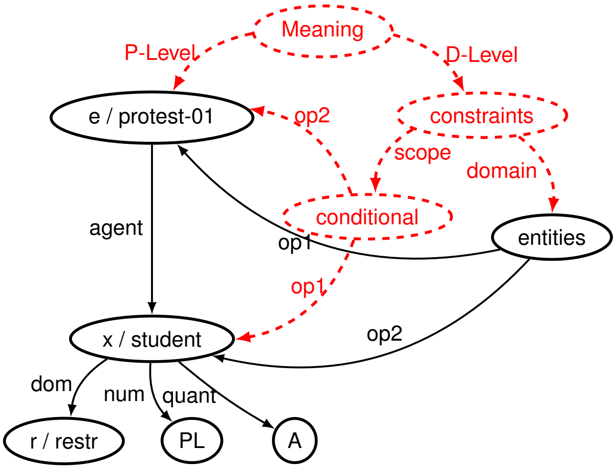

Figure 2: Enriched AMR with Quantification

The encoding of the contextual constraints, represented on the graph, may be treated by the indexing mechanism referred to in Bos (2020), but we have taken a different option.

4. Validation with Further Illustrations

4.1. Overview

This section informally validates the proposed two-block reference semantics with further illustrations. Various discourse constraints, such as scope involving quantification or presuppositional existence associated with definite descriptions or proper names, can be represented by incorporating Bos (2020)'s contextual constraints into the *discourse level*. These contextual constraints allow AMR to avoid encoding some information redundantly. With conditional indexing for universal quantification, Bos even argues that the traditional `:quant all` marking can be dispensed with. In the present proposal, however, it is better to retain overtly expressed information, such as determinacy, rather than relying on contextual information alone.

4.2. Definite Descriptions: Negation and Presupposition

Definite descriptions, such as *my cat* and *the rats*, are known to trigger existential presupposition. Consider how definite description phrases interact with negation in a sentence.

(17) :: Sentence: My cat didn't chase out all the rats from the house.
:: Interpretation: I had a cat, and there were rats in my house, but it is not the case that the cat chased all of the rats out of the house.
:: DCT 2026-02-11

Sentence (17) contains three definite descriptions, *cat*, *rats*, and *house*. Encoding each of them with `:determinacy (d / def)`, as in QuantML, presupposes their existence.

(18)
```
(e / chase-01
   :polarity -
   :agent (c / cat
            :poss (m / me)
            :determinacy
              (d1 / def))
   :theme (r / rat
            :determinacy
              (d2 / def))
   :source (h / house
            :determinacy
              (d3 / def))
   :goal (o / out))
(cn / constraints
   :presuppose (ex / exists
                  :agent c
                  :theme r
                  :source h)
   :scope (w1 / wide
              :op1 -
              :op2 e)
   :scope (w2 / wide
              :op1 r
              :op2 e))
```

Annotation (18) yields Logical Form (19):

(19) Logical Form:
a. Block 1:
$\{e, c, m, r, h, o, d_1, d_2, d_3\}$
$[instance(e, chase_{01}),$
$instance(c, cat), instance(d_1, def),$
$instance(m, me),$
$instance(r, rat), instance(d_2, def),$
$instance(h, house), instance(d_3, def),$
$instance(o, out),$
$polarity(e, -),$
$arg_0(e, c), arg_1(e, r),$
$source(e, h), goal(e, o)]$
b. Block 2, referring to Block 1:

$\neg\forall y[y{=}r \rightarrow$
$[instance(e, chase_{01}), agent(e, c),$
$theme(e, y), source(e, h), goal(e, o)]]$

The first part of Logical Form (19) states that the described situation introduces an event of chasing and the discourse entities corresponding to the cat, the rats, and the house. The second part states the weak reading of the negated universal: it is not the case that every relevant rat was chased out of the house by the cat. If the negation operator were placed elsewhere, the truth conditions would change while the existential presuppositions remain in force.

4.3. Subevents

The AMR Guidelines (Banarescu et al., 2019) list `:subevent` and `:part` among the non-core relations. These relations can help encode the internal structure of complex events such as sharing something edible or a house, the Olympic Games with many component events, or wars with dispersed battles.

Consider a case of two women sharing a piece of pizza, taken from Bunt and Lee (2025). The AMR annotation is simplified by using the relation `:subevent`.

(20) :: Sentence: Two women shared a pizza.
:: Interpretation: There were two women and a pizza, and they shared it, each eating part of it.
:: CDT-2026-01-28:T20:10

```
(s / share-01
   :agent (w / woman
            :quant 2)
   :theme (p / pizza
            :quant 1)
   :manner collective
   :subevent
     (e / eat-01
       :agent w
       :theme (p1 / pizza
                  :part-of p)
       :manner individual))
(c / constraints
   :scope (w1 / wide
             :op1 p
             :op2 w)
   :scope (w2 / wide
             :op1 p1
             :op2 e))
```

In the intended scenario, sharing the pizza is modeled as including eating subevents. The eating event is therefore annotated as a subevent of the sharing. Logical Form (21) below is derived from Annotation (20).

(21) Logical Form:
a. Block 1: Preamble
$\{s, w, p, e, p1\}$
$[instance(s, share_{01}),$
$instance(w, woman), quant(w, 2),$
$instance(p, pizza), quant(p, 1),$
$agent(s, w), theme(s, p),$
$manner(s, collective),$
$instance(e, eat_{01}),$
$subevent(e, s), agent(e, w),$
$instance(p1, pizza), part\_of(p1, p),$
$theme(e, p1), manner(e, individual)]$
b. Block 2, referring to Block 1:
$\forall e_1[e_1 = e \rightarrow [instance(e_1, eat_{01}),$
$agent(e_1, w), theme(e_1, p1)]]$

Logical Form (b) again refers to the preamble for the variables supplied by Block 1.

4.4. Distributivity and Coreference

Here is another example from QuantML (ISO (2025), Example 5) involving distributivity and coreference.

(22) :: Sentence: The men had a beer, and they carried the piano upstairs.
:: Interpretation: There were some men and a piano. Under the intended reading, each man drank a beer, and the same group then carried the piano upstairs together.
:: id: k006, DCT 2025-01-04

```
:predicate-level
 (d / drink-01
    :agent (m / man
             :num PL)
    :theme (b / beer
             :quant 1)
    :manner individual)
 (c / carry-01
    :agent (m1 / man
             :num PL)
    :theme (p / piano
             :quant 1)
    :goal (u / upstairs)
    :manner collective)
:discourse-level
 (cn / constraints
     :presupposition
         existential
     :temporal (bf / before
                  :op1 d
                  :op2 c)
     :scope (sc / conditional
                :op1 m
                :op2 d)
     :coreference
         (id / identical
            :op1 m :op2 m1))
```

The individual distributivity of the drinking event triggers the scope constraint over the men, while the carrying event is marked as collective. The coreference relation states that the men who drank are the same men who later carried the piano.

(23) Logical Forms:
a. $\{d, m, b, c, m1, p, u\}$
$[instance(d, drink_{01}),$
$instance(m, man), num(m) \geq 2,$
$instance(b, beer), quant(b, 1),$
$instance(c, carry_{01}),$
$instance(m1, man), num(m1) \geq 2,$
$instance(p, piano), quant(p, 1),$
$instance(u, upstairs),$
$agent(d, m), theme(d, b),$
$manner(d, individual),$
$agent(c, m1), theme(c, p),$
$goal(c, u),$
$manner(c, collective)]$
b.
$\forall y[y{=}m \rightarrow$
$\exists e_1[e_1 \sqsubseteq d, agent(e_1, y),$
$\tau(d) \prec \tau(c), m{=}m_1)],$
where τ maps an event to the time of its occurrence (event time).

This logical form makes the distributive reading explicit for the drinking event while keeping the collective carrying event and the coreference relation at the discourse level.

5. Conclusion

Quantificational meaning is pervasive in text, but it remains difficult to express in standard AMR under the common "flat conjunction" interpretation. In particular, universal quantification and distributive readings require scope-taking structure and variable binding that cannot be recovered from a mere list of conjuncts. This paper presented a conservative extension of AMR that makes quantificational commitments explicit while preserving AMR's practical graph backbone and remaining compatible with the two-layer architecture of UMR..

The paper makes five concrete contributions. First, we proposed an AMR-compatible representation in which a sentence-level predicate-argument graph (P-level) is paired with a document/discourse component (D-level) that records contextual constraints relevant to quantifier interpretation. The two levels are explicitly *linked* by shared identifiers, allowing discourse constraints (e.g., scope) to target the same event and participant nodes introduced at the predicate level.

Second, we showed how key ISO 24617-12:2025 QuantML dimensions (including *domain restriction*, *determinacy*, *distributivity*, and *involvement*) can be represented as local, meaning-bearing annotations on AMR nodes/edges in PENMAN, avoiding the verbosity of the XML concrete syntax while preserving the underlying annotation content.

Third, building on the spirit of contextual constraints (e.g., AMR$^+$), we defined D-level constraint structures (e.g., conditional scope constraints, presuppositional existence constraints, and cross-sentential linking) that encode scopal dependencies and discourse-sensitive conditions without introducing a separate layer of explicit index bookkeeping inside the graph.

Fourth, we provided a systematic mapping from enriched graphs to a two-block logical form: (i) a *minimal model* block that records the described situation (events and discourse referents contributed by the AMR predicate structure), and (ii) a *constraint* block that states quantificational and discourse-sensitive requirements (e.g., universal/conditional structure, presuppositions, temporal constraints) while allowing controlled reference back to Block 1.

Finally, through detailed examples, we illustrated how the proposal captures (i) restricted universals as conditional scope constraints, (ii) interactions between definiteness, negation, and presuppositional existence, and (iii) distributive interpretations that require event-structural enrichment (e.g., event-part quantification and subevent structure) while maintaining a single graph-based interface.

Overall, the proposal aimed to make quantification representable and checkable in AMR-style annotation. QuantML feature content remains locally readable at the predicate level, and scope-taking and discourse-sensitive commitments are stated explicitly at the discourse level to support interpretation and downstream reasoning.

Next steps include (i) expanding and systematizing the discourse-level constraint inventory (including additional triggers from QuantML features), (ii) developing annotation guidelines and consistency criteria for the enriched representation, and (iii) conducting empirical evaluation, including annotation studies and experiments on parsing and downstream tasks that are sensitive to quantificational distinctions (e.g., inference under scope and distributivity).

6. Acknowledgements

We thank the three anonymous reviewers for their detailed and constructive comments, which have substantially improved the paper. We have addressed most of their suggestions in the present revision. We also owed a great deal to three scholars: Harry Bunt, Johan Bos, and James Pustejovsky, for their original ideas, which we tried to incorporate into this seminal work. A few broader issues,

mainly concerning semantic interpretation, are beyond the scope of this paper and are left for future work.

7. Bibliographical References

Emmon Bach, Eloise Jelinet, Angelica Kratzer, and Barbara H. Partee, editors. 1995. *Quantification in Natural Language*. Kluwer Academic Pubishers, Dordrecht.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2019. *Abstract Meaning Representation (AMR) 1.2.6 Specification*. The DMR Working Group.

Jon Barwise. 1981. Scenes and other situations. *The Journal of Philosophy*, 78(3):369–397.

Patrick Blackbun and Johan Bos. 2003. Computational semantics. *Theoria*, 18(1):27–45.

Patrick Blackburn and Johan Bos. 2004. *Represenging and Inference for Natural Language: A First Course in Computational Semantics*. Center for the Study of Language and Information, Stanford, CA.

Johan Bos. 2003. Exploring model building for natural language understanding. In *ICoS-4. Inference in Computational Semantics. Workshop Proceedings*, pages 41–55.

Johan Bos. 2016. Squib: Expressive power of abstract meaning representations. *Computational Linguistics*, 42(3):527–535.

Johan Bos. 2020. Separating argument structure from logical structure in AMR. In *Proceedings of the 2nd Interantional Conference on Designing Meaning Representations*, pages 13–20, Barcelona, Spain (online), December 13, 2020.

Harry Bunt. 2010. A methodology for designing semantic annotation languages exploiting syntactic-semantic iso-morphisms. In *Proceedings of ICGL 2010, the Second International Conference on Global Interoperability for Language Resources*, pages 29–45, City University of Hong Kong.

Harry Bunt, Maxime Amblard, Johan Bos, Karën Fort, Bruno Guillaume, Philippe de Groote, Chuyuan Li, Pierre Ludmann, Michel Musiol, Siyana Pavlova, Guy Perrier, and Sylvain Pogodalla. 2022. Quantification annotation in ISO 24617-12, second draft. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3407–3416, Marseille, France. European Language Resources Association.

Harry Bunt and Kiyong Lee. 2025. The representation of QuantML annotations in UMR – an exploration. In *Proceedings of the 21th Joint ACL–ISO Workshop on Interoperable Semantic Annotation (ISA-21)*, Heinrich Heine University Düsseldorf, Germany. IWCS 2025.

Harry Bunt, James Pustejovsky, and Kiyong Lee. 2018. Towards an ISO standard for the annotation of quantification. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*, pages 1787–17949, Miyazaki, Japan. ELRA.

ISO. 2016. *ISO 24617-6:2016, Language resource management – Semantic annotation framework (SemAF) – Part 6: Principles of semantic annotation (SemAF Principles)*. The International Organization for Standardization, Geneva. Project leader: Harry Bunt.

ISO. 2025. *ISO 24617-12:2025 Language resource management – Semantic annotation framework (SemAF) – Part 12: Quantification*. The International Organization for Standardization, Geneva. Project leader: Harry Bunt.

Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic: An Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and DRT*. Kluwer, Dordrecht.

Kevin Knight, Bianca Badarau, Laura Baranescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O'Gorman, and Nathan Schneider. 2020. *Abstract Meaning Representation Release 3.0*. Linguistic Data Consortium, Philadelpia, PA. Techinical Report LDC2020T02.

Kiyong Lee, Harry Bunt, James Pustejovsky, Alex C. Fang, and Chongwon Park. 2025. Representing ISO-annotated dynamic information in UMR. In *Proceedings of the Sixth International Workshop on Designing Meaning Representations*, pages 49–58, Prague, Czechia. Association for Computational Lingustics.

Richard Montague. 1974. The proper treatment of quantification. In *Formal Philosophy: Selected Papers of Richard Montague*, New Haven and London. Yale University Press.

James Pustejovsky, Nianwen Xue, and Kenneth Lai. 2019. Modeling quantification and scope in abstract meaning representations. In *Proceedings of the First International Workshop on*

Designing Meaning Representation, pages 28–33. Association for Computational Linguistics. Florence, Italy, August 1, 2019.

Working Group UMR. 2022. *Uniform Meaning Representation (UMR) 0.9 Specification*. UMR Working Group for Guidelines.