

A semantic annotation scheme for quantification

Harry Bunt

Department of Cognitive Science and Artificial Intelligence

Tilburg University, The Netherlands

harry.bunt@uvt.nl

Abstract

This paper describes in brief an annotation scheme called 'QuantML' which was proposed last December to the International Organisation for Standardisation (ISO) as a starting point for developing a standard for interoperable rooted in the theory of generalised quantifiers, neo-Davidsonian semantics, and DRT, covers a wide range of aspects of quantification. The scheme consists of (1) an abstract syntax which defines 'annotation structures' as triples and other set-theoretic constructs; (b) a compositional semantics of annotation structures; (3) an XML representation of annotation structures.

1 Introduction

Quantification is widespread in spoken and written language; it can be found in nearly every sentence, since it occurs whenever a predicate is applied to one or more argument sets (rather than single arguments). This commonly happens in two of the linguistically most prominent units: clauses and noun phrases. In clauses it happens when a verb is combined with sets of arguments, as in "*Last year the American car manufacturers produced more than 12 million vehicles*". In noun phrases it happens when a head noun is subject to modification, as in "*I'm carrying some heavy books*".

A widely held view is that the quantifiers of natural language are not determiners like "*some*" and "*all*", in spite of their superficial similarity with the quantifiers of formal logic, but rather noun phrases, like "*more than 12 million vehicles*" and "*some heavy books*". Other types of quantifiers can also be found, such as adverbs for temporal or spatial quantification ("*always*", "*nowhere*", "*sometimes*"), but these are of minor importance compared to noun phrases. Hobbs and Shieber (1987) have argued that the sentence "*Some representatives of every department in most companies saw a few samples of every product*", containing five noun phrases, has 42 readings, corresponding to equally many linguistically valid alternative scopings of the five quantifiers (of the 120 mathematically possible permutations). Bunt and Muskens (1999) show that, when other ambiguity types are taken into account, such as those of a quantification's distributivity, the number of possible readings of an ordinary written sentence may run into the thousands.

Quantification is the main source of structural ambiguity in natural language; therefore applications of natural language processing for which semantic information is important, such as information extraction, and question answering, need effective ways of interpreting quantification expressions. This calls for a flexible and interoperable way of indicating aspects of quantification.

This paper presents an annotation scheme, called 'QuantML', where a range of aspects of natural language quantification are captured by a relatively small number of features, distributed over the components of annotation structures in a way that allows a compositional semantic interpretation. The QuantML scheme is based on a number of preliminary studies including Bunt (2017), Bunt et al. (2018), and Bunt (2018), and has recently been accepted as the basis for developing Part 12 of the ISO Semantic annotation framework (ISO 24617).

2 Related work

ISO standard 24617-1 for the annotation of time and events, commonly known as ‘ISO-TimeML’, has certain provisions for dealing with time-related quantification. For example, the temporal quantifier “daily” is represented as follows, where “PID” stands for “period of one day”:

(1) <TIMEX3 xml:id=“t5” target=“#token0” type=“SET” value=“PID” quant=“EVERY”/>

In ISO-TimeML @quant is one of the attributes of temporal entities, used to indicate that the entity is involved in a quantification. The limitations of this approach for annotating temporal quantification have been discussed by Bunt & Pustejovsky (2010).

ISO-Space (ISO 24617-7) uses the @quant attribute as well, 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:

(2) a. There’s a lamp_{se1} on_{ss1} every desk_{se2}.

b. <spatialEntity id=“se1” target=“#token2” pred=“lamp” form=“nom” countable=“true” quant=“1” scopes=“0”/>
<spatialEntity id=“se2” target=“#token5” pred=“desk” form=“nom” countable=“true” quant=“every” scopes=“#se1”/>
<spatialSignal id=“ss1” target=“#token3” type=“dirTop” />
<qsLink id=“qsl1” relType=“EC” figure=“#se1” ground=“#se2” trigger=“#ss1”/>
<oLink id=“ol1” relType=“above” figure=“#se1” ground=“#se2” trigger=“#ss1” frameType=“intrinsic” referencePt=“#se2” projective=“false” />

From a semantic point of view, this use of the @scopes attribute is not very satisfactory 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 (2) does not have a compositional semantics.

The Parallel Meaning Bank (PMB, Abzianidze et al., 2017), building upon the Groningen Meaning Bank (GMB, Bos et al., 2017), is a corpus of semantically annotated sentences and texts in English, German, Dutch and Italian in raw and tokenised format with formal meaning representations. The PMB aims to provide fine-grained meaning representations in DRT for the most likely interpretation of a sentence, with a minimal use of underspecification. The GMB and the PMB are very useful resources for semantic study, but this work is somewhat different from the usual kind of annotation work, where semantic features are associated with small stretches of text or speech.

3 Granularity in quantification annotation

3.1 Ambiguity and lack of specificity

The multiplicity of possible readings of quantifications forms a challenge for language understanding systems, but hardly for humans, who are mostly not aware of the ambiguities. Human annotators who are not trained linguists or logicians likewise tend not to see all the possible readings of quantifications. Automatic annotation processes run into the same problems as language understanding systems, having a lack of general world knowledge and situation-specific context information. Both for manual and for automatic annotation it is therefore of practical importance to not be forced to make more specific choices than the available information and skills justify. On the other hand, it should of course be possible for a skilled annotator to make precise annotations if sufficiently detailed information is available. A useful annotation scheme should thus allow specifications with varying degrees of granularity.

The ambiguity challenge that quantification poses for automatic language understanding has led to the development of underspecification techniques in computational semantics, in particular for underspecified representation of quantifier scope (e.g. Alshawi, 1992; Bos, 1995; Reyle, 1993, Willis & Man-

andhar, 2001). In the same vein, QuantML allows the annotation of quantifier scope to be underspecified by making the specification of scope relations between a pair of quantifiers optional.

Scope is not the only source of quantifier ambiguity; ambiguities in the distributivity and individuation of quantifiers make the ambiguity problem even more dramatic than has generally been assumed in the literature, due to issues concerning precision and individuation which are considered next.

3.2 Precision and distributivity

Ambiguity in natural language quantification is mostly considered in terms of the number of logically precise interpretations. But natural language expressions are sometimes not meant to be interpreted with logical precision. This is in particular the case for quantifier distributivity. Consider the following example:

- (3) The men carried all the boxes upstairs.

In the event(s) described by this sentence it is not necessarily the case that all the carrying was done either collectively or individually; the sentence could for instance describe a set of events in which the men collectively carried the heaviest boxes, and individually the lighter ones. This means that the distributivity of the quantification over the set of men is neither collective nor individual (and the same is true for the quantification over the boxes); the term ‘unspecific’ has been introduced for this distributivity (Bunt, 1985). This interpretation can be represented in second-order predicate logic as shown in (4), where following Kamp & Reyle (1993)¹ the notation X^* is used to designate the set consisting of the members and subsets of X , and moreover the subscript notation P_0 to designate the characteristic predicate of the reference domain of a quantifier,² which is a contextually determined part of the quantifier’s source domain (as determined by an NP head), characterized by the predicate P .

- (4) $\forall x [\text{box}_0(x) \rightarrow \exists y \exists e [\text{man}_0^*(y) \wedge \text{carry-up}(e) \wedge \text{agent}(e,y) \wedge \exists z [\text{box}_0^*(z) \wedge [x=z \vee x \in z] \wedge \text{theme}(e,z)]]]$

This representation says that for every box in a given reference domain there is a carry-event in which a contextually distinguished man or group of men carried it upstairs or carried a set of boxes upstairs that contains it.

A quantification with unspecific distributivity has both individual and collective participation as special cases, so ‘unspecific’ could be used to avoid having to choose a more specific distributivity. In the case of (3), ‘unspecific’ is the correct distributivity to assign. In cases where it is difficult to decide on the distributivity, ‘unspecific’ could be useful as a coarse-grained default value in annotations.³

3.3 Individuation

The ‘individuation’ of a quantification is another source of ambiguity, as illustrated in (5):

- (5) a. I see no chicken in the garden.
b. I see no chicken in the stew.

The count/mass distinction is often characterized semantically in terms of ‘individuation’: “*To learn ‘apple’ ... we must learn how much counts as an apple. (...) Such terms possess built-in modes (...) of dividing their reference (...) Consider ‘shoe’, ‘pair of shoes’, and ‘footwear’: (...) two of the them divide their reference differently, and the third not at all.*” (Quine, 1960). In other words, count nouns have a domain of reference made up of individuals, while that of a mass noun is made up of entities (often called ‘quantities’) with mereological part-whole relations.

¹Kamp & Reyle 1983, Section 4.2.2

²Also known as ‘context set’, Westerståhl, 1985.

³See also Schwertel, 2005 for a discussion of vagueness and lack of specificity in quantification.

Quantifiers expressed by an NP with a count head noun may be ambiguous in a different way, as illustrated by (6a). This sentence could for example describe a series of events where last Monday Mario had a pizza, on Wednesday he had another pizza plus a few slices, and on Friday he had the slices remaining from Wednesday. Pizzas are a domain where it is common to consider parts of individuals, like in many domains related to food and drink. For other domains this may be less common, but in principle every physical object has parts, and many abstract objects as well. When interpreting an NP that describes domain involvement or domain size in terms of a non-integer number of individuals, this is clearly necessary. The interpretation of sentence (6a) as describing a set of events in which Mario has eaten some pieces of pizza, adding up to a total of three pizzas, can be represented by (6b), where P^+ designates the property of being a part of an individual that has the property P ,⁴ and Σ designates the joining together of parts of an individual.⁵

(6) a. Mario had three pizzas last week.

b. $\exists Y (\forall y (y \in Y \rightarrow (\text{pizza}^+(y) \wedge \exists e (\text{eat}(e) \wedge \text{agent}(e, \text{Mario}) \wedge \text{theme}(e, y)))) \wedge |\Sigma Y|^{\text{pizza}}=3)$

In model-theoretic semantics it is commonly assumed that individuals are atomic concepts, but for examples like the above we must assume an ontology where individuals have parts (in the mereological sense). This part-whole relation has the same logical properties as the corresponding relation for mass nouns.⁶ A quantification where parts of individuals should be taken into account will be said to have the individuation “*count/parts*”.

Individuation and distributivity are distinct aspects of quantification; elements from a domain with count/parts individuation can for example participate collectively in a quantification, as in the report from a weight-lifting contest that “*Tarzani lifted 27.5 pizzas*” (see (15) below), if Tarzani lifted a pile of 25 whole pizzas topped by a short stack of five pizza halves.

4 Methodology

4.1 Theoretical background

The theory of generalized quantifiers (GQT) has been successful in describing and understanding many aspects of natural language quantification. Together with neo-Davidsonian event semantics and Discourse Representation Theory (DRT), GQT forms the theoretical basis of the approach to quantification annotation taken in this paper.

GQT exploits the fact that quantification in natural language differs from that in formal logic in that logical quantifiers like “for all x ” and “there exists an x ” range over all the individual objects in a given universe of discourse, whereas quantifying expressions in natural language like “*all the students*”, “*an essay*”, “*some coffee*”, indicate a restricted domain that the quantification applies to. GQT therefore views noun phrases as the quantifiers of natural language (rather than determiners) (Barwise and Cooper, 1981; see e.g. also van Benthem and ter Meulen, 1985 and Szabolcsi, 2010). This view generalizes to natural language quantifiers like “*two books*”, “*less than three weeks*”, “*thirty tons of peanut butter*”. Semantically, generalized quantifiers are viewed as expressing properties of sets of individuals; for example, the quantifier “*more than three essays*” is interpreted as the property of being a set that contains more than three essays.

⁴Formally, $P^+(x) \leftrightarrow \exists y (P(y) \wedge x \sqsubseteq y)$.

⁵Expressing the size of a collection of pizza-parts in terms of number of pizzas assumes (1) a way of measuring the size of pizzas and pizza-parts, and (2) a standard size. The notation $|\cdot|^D$ designates such a measure function for the domain D . Note that, for example, eight quarters of three different pizzas together have a size of 2 pizzas, even though it may not be possible to physically join the parts to form two well-formed pizzas. The mereological sum ‘ Σ ’ in (6b) is used to avoid counting the size of overlapping parts more than once.

⁶The most important properties are: (1) a part of a part of an individual is again a part of that individual: $(x \sqsubseteq a \wedge y \sqsubseteq x) \rightarrow y \sqsubseteq a$; (2) the mereological sum of parts of an individual is another part of that individual: $(x \sqsubseteq a \wedge y \sqsubseteq a) \rightarrow \Sigma(x, y) \sqsubseteq a$; (3) every object is part of itself: $x \sqsubseteq x$. See Bunt (1985) for a detailed discussion of the logic of mereological part-whole relations, including a discussion of Lesniewski’s mereology and Leonard & Goodman’s ‘calculus of individuals’.

Davidson (1967) proposed to treat events as individual objects, facilitating the semantic interpretation of adverbs, like “quickly”, “passionately”, and adverbial quantifying expressions such as “everywhere”, “never” and “at least three times”. Following Parsons (1990), this event-based semantics can be expressed in semantic representations by means of one-place predicates applied to existentially quantified event variables, and two-place predicates to indicate the semantic roles of the participants in an event. This ‘neo-Davidsonian’ approach has been adopted in the ISO annotation standards 24617-1 (Time and events), 24617-4 (Semantic roles), 24617-7 (Spatial information), and 24617-8 (Discourse relations). Champollion (2015) has shown that GQT and neo-Davidsonian semantics can be combined successfully. Still, natural language quantification is a semantically extremely complex set of phenomena, and especially the interpretation of plural noun phrases presents certain theoretical challenges for GQT (see e.g. Schwertel, 2005), some of which have been successfully been approached in DRT (Kamp and Reyle, 1993), which has other limitations. Luckily, providing a semantics for quantification *annotations* is less challenging than providing a semantics for natural language expressions involving quantifications.

Several of the ISO semantic annotation standards use DRT’s Discourse Representation Structures (DRSs) for defining a semantics of annotation structures. QuantML follows suit, combining ideas from GQT, neo-Davidsonian semantics, and DRSs in the semantics of its annotation structures.

4.2 Annotation scheme architecture

The annotation scheme outlined in this paper has been designed according to the ISO principles of semantic annotation (ISO standard 24617-6). This means that the scheme has a three-part definition consisting of (1) an abstract syntax that specifies the possible *annotation structures* at a conceptual level as set-theoretical constructs, such as pairs and triples of concepts; (2) a semantics that specifies the meaning of the annotation structures defined by the abstract syntax; (3) a concrete syntax, that specifies a representation format for annotation structures using XML expressions. Defining the semantics at the level of the abstract syntax puts the focus of an annotation standard at the conceptual level, rather than that of representation formats. Alternative representation formats may be defined with guaranteed interoperability (Bunt et al., 2018). Annotators (human or automatic) deal with concrete representations only, but they can rely on the existence of an underlying abstract syntax and semantics.

Example (7) shows the QuantML annotation structure (in a slightly simplified form), XML representation, and semantics for the collective reading of a simple sentence. Besides the usual box notation of DRT also a string notation will be used, which is shown in (7e). Capital variables are used to designate non-empty sets.

(7) a. Two thousand students protested

Markables: m1 = Two thousand students; m2 = students; m3 = protested

b. QuantML annotation structure:

$\langle \langle m3, \langle \text{protest} \rangle \rangle, \{ \langle m1, \langle \text{student}, \lambda z. |z|=2000, \text{indef} \rangle \rangle \}, \text{agent, collective} \rangle, \{ \}$


c. Annotation representation:

`<entity xml:id="x1" target="#m1" pred="student" involvement="2000"/>`

`<entity xml:id="e1" target="#m3" pred="protest"/>`

`<participantLink event="#e1" participant="#x1" semRole="agent" distr="collective"/>`

d. Semantics:



e. $[X, E \mid |X|=2000, [x \in X \rightarrow \text{student}(x)], [e \in E \rightarrow [\text{protest}(e), \text{agent}(e,X)]]]$

It may be noted that the annotation semantics in (7d,e) is structurally the same as the DRS that Kamp and Reyle propose for the collective reading of the sentence “*Three lawyers hired a new secretary*” (Kamp and Reyle 1993, p. 327). For the individual reading of the sentence (7a), where the students act in individual protest-events (e.g. writing personal letters of protest), the annotation structure and its XML representation would both have ‘individual’ instead of ‘collective’ (and narrow event scope, by default), and the DRS interpretation would be as in (8):

$$(8) [X, E \mid |X|=2000, [x \in X \rightarrow \text{student}(x)], [e \in E \rightarrow \text{protest}(e)], [x \in X \rightarrow [e \mid e \in E, \text{agent}(e,x)]]]$$

Note also that the discourse referent X in these DRSs stands for the set of entities that participate in the protest events, which corresponds to the set of entities (or the property) that in a classical linguistic analysis is denoted by the VP. The DRS thus has a condition of the form $[x \in X \rightarrow \dots\theta(e,x)]$ for the individual reading and a condition with $\theta(e,X)$ for the collective reading. The two conditions $|X|=2000$, $[x \in X \rightarrow \text{student}(x)]$ together reflect the GQT interpretation of the subject NP.

5 QuantML specification

5.1 Abstract syntax

The structures defined by the abstract syntax are n-tuples of elements that are either basic concepts, taken from a store called the ‘conceptual inventory’, or, recursively, such n-tuples. Two types of structure are distinguished: entity structures and link structures. An entity structure contains semantic information about a segment of primary data and is formally a pair $\langle m, s \rangle$ consisting of a markable, which refers to a segment of primary data, and certain semantic information. A link structure contains information about the way two or more segments of primary data are semantically related.

QuantML conceptual inventory:

- unary predicates that characterize source domains (such as ‘book’, ‘student’, and ‘water’) or event domains (such as ‘lift’, ‘carry’, ‘drink’), or that correspond to adjectives or to prepositions;
- binary predicates that correspond to semantic roles (agent, theme, instrument,...); for this purpose, the semantic roles defined in ISO 24617-4 (Semantic roles) are used;
- numerical predicates for specifying reference domain involvement, reference domain size, the size of certain parts of a reference domain, or the number of repetitions or frequency of recurrence in event structures;
- predicates for specifying proportional reference domain involvement, such as ‘all’, ‘some’, ‘most’;
- parameters for specifying definiteness: ‘definite’ and ‘indefinite’; domain individuation: ‘count’, ‘mass’, and ‘count/parts’; and distributivity: ‘collective’, ‘individual’, ‘homogeneous’, ‘group’, ‘unspecific’;
- basic units of measurement, such as ‘meter’, ‘kilogram’, ‘litre’, and the operators ‘division’ and ‘multiplication’ for forming complex units;
- the polarity values ‘positive’ and ‘negative’;
- parameters for specifying event scope: ‘wide’ and ‘narrow’, and for specifying whether scope inversion occurs: ‘inverse’ or ‘linear’ (default value).
- ordering relations for specifying the relative scopes of quantifiers over sets of participants: ‘wider’, ‘equal’, ‘dual’, and ‘unscoped’.

QuantML has three kinds of entity structures: (1) for events; (2) for participants; (3) for restrictions on sets of participants. A quantified set of participants is characterized by the following properties:

1. the source domain, from which the participants are drawn;
2. the reference domain, typically a subset of the source domain;
3. the individuation of the reference domain;
4. the quantitative (absolute or proportional) involvement of the reference domain;
5. the size of the reference domain, or of groups, subsets, or parts of the reference domain involved in the quantified predication (optional).

The entity structure $\langle m, s \rangle$ for a set of participants thus contains a triple $s = \langle \langle D, v \rangle, q, d \rangle$ with D = characteristic domain predicate, v = individuation, q = reference domain involvement, and d = definiteness, with possibly an additional size specification. The domain component is more complex when the restrictor of an NP contains one or more head noun modifiers and/or multiple, conjoined heads (see Bunt 2018 for details).

Entity structures for sets of events are simpler than those for participants; they contain just a predicate that characterizes a domain of events, and if applicable the cardinality of a set of repeated events or the frequency of recurring events.

Two kinds of link structure are defined: participation structures, which link participants to events, and scope link structures. Participation structures specify (1) a set of events; (2) a set of participants; (3) a semantic role; (4) the distributivity of the participation; (5) the relative scope of the event quantification. Scope link structures specify the relative scope of two participant entity structures.

Annotation structures for quantification are associated mostly with clauses and their constituent NPs and verbs. The annotation structure for a clause is a quadruple consisting of an event structure, a set of participant structures, a set of participation link structures, and a (possibly empty) set of scope link structures. In a complete clause annotation structure all participant entity structures are linked to the verb's event entity structure, and all the relative scopes of all participant entity structures are specified.

5.2 Concrete syntax

A concrete syntax is specified here in the form of an XML representation of annotation structures. For each type of entity structure, defined by the abstract syntax, a corresponding XML element is defined; each of these elements has an attribute `@xml:id` whose value is a unique identifier, and an attribute `@target`, whose value anchors the annotation in the primary data, having a markable as value (or a sequence of markables). In addition, these elements have the following attributes:

`<entity>`: `@domain`, `@involvement`, `@definiteness` and `@size` (optional);
`<event>`: `@pred` (event type), `@number` (optional), and `@frequency` (optional);
`<qDomain>`: `@source` (with multiple values in the case of a conjunctive head) and `@restrictions`;
`<sourceDomain>`: `@pred`, `@individuation`;
`<adjMod>`: `@pred`, `@distr`, and `@restrictions` (optional);
`<nnMod>`: `@pred`, `@distr`, and `@restrictions` (optional);
`<ppMod>`: `@pRel`, `@pEntity`, `@distr`, `@linking`;
`<relClause>`: `@semRole`, `@clause`, `@distr`, `@linking`;
`<amount>`: `@num`, `@unit`;
`<complexUnit>`: `@unit1`, `@operation`, `@unit2`.

For each of the two types of link structure defined by the abstract syntax, an XML element is defined:

`<participantLink>` with the attributes `@event`, `@participant`, `@semRole`, `@distr`, `@evScope` (default value: narrow) and `@polarity` (default value: positive);
`<scopeLink>` with the attributes `@arg1`, `@arg2`, `@scopeRel`.

5.3 Semantics

The QuantML semantics specifies a recursive interpretation function I_Q that translates annotation structures into DRSS in a compositional way: the interpretation of an annotation structure is obtained by combining the interpretations of its component entity structures and participation link structures, in a way that is determined by scope link structures (if any). A full specification of the QuantML semantics would go beyond the scope of this paper; the reader is referred to Bunt (2018, Appendix C). Here we outline the overall approach and present some interesting parts of the definition of I_Q .

The QuantML interpretation function translates every participant entity structure, event entity structure, and participation link structure into a DRS and combines these. Consider the example in (7). The entity structures for “*Two thousand students*”, and “*protested*” are translated into the DRSs shown in (10). For the participant entity structure this is achieved by applying an instance of clause (9a) in the I_Q definition, which interprets entity structures with source domain D , individuation v , involvement q , and definiteness $indef$. The interpretation q' of domain involvement specification q is defined in (9b-c), and that of the domain specification in (9d-e).

- (9) a. $I_Q(\langle m1, \langle \langle m2, D \rangle, v \rangle, q, indef \rangle) = [X \mid q'(X), [x \in X \rightarrow D'(x)]]$
 b. $q' = I_Q(q) \circ F_Q(v)$
 c. $F_Q(v): F_Q(count) = \lambda X. |X|$; $F_Q(mass) = F_Q(count/parts) = \lambda X. |\Sigma X|$
 d. $D' = I_Q(\langle D, v \rangle)$
 e. $I_Q(\langle D, v \rangle): I_Q(\langle D, count \rangle) = I_Q(\langle D, mass \rangle) = I_Q(D)$; $I_Q(\langle D, count/parts \rangle) = I_Q(D)^+$

- (10) a. $I_Q(\langle m1, \langle \langle m2, student \rangle, count \rangle, \lambda z. |z| = 2000, indef \rangle) = [X \mid |X|=2000, [x \in X \rightarrow student(x)]]$
 b. $I_Q(\langle m3, \langle protest \rangle) = [E \mid [e \in E \rightarrow protest(e)]]$

The DRS in (10a) says that there exists a set with the property of containing two thousand students, reflecting the GQT approach to NP interpretation. The DRS in (10b) together with (12) illustrates the adoption of neo-Davidsonian event semantics.

The participation link structure has in this example the form $\langle \varepsilon_E, \{ \varepsilon_{P1} \}, R, d, \sigma \rangle$, where ε_E and ε_{P1} are the participant and event entity structures that are linked in the Agent role ($R = Agent$), with $d =$ collective, and σ (event scope) = narrow. The semantic interpretation of such a structure is defined as follows, where ‘ \cup ’ designates the familiar merge operation for DRSs:

- (11) $I_Q(\langle \varepsilon_E, \{ \varepsilon_{P1} \}, R, d, \sigma \rangle) = I_Q(\varepsilon_{P1}) \cup (I_Q(\varepsilon_E) \cup I_Q(R, d, \sigma))$

Triples like $\langle R, d, \sigma \rangle$ are interpreted as shown in (12):

- (12) a. $I_Q(R, individual, narrow) = [X, E \mid x \in X \rightarrow [e \mid E \in E, agent(e, x)]$
 b. $I_Q(R, individual, wide) = [X, E \mid x \in X \rightarrow [e \mid E \in E, agent(e, x)]$
 c. $I_Q(R, collective, \sigma) = [X, E \mid x \in X \rightarrow [e \mid E \in E, R(e, X)]$

Applying rule (11) to the right-hand sides of (10) and (12c), with the values for R, d and σ substituted, gives the desired result shown in (7d,e):

$$[X, E \mid |X|=2000, [x \in X \rightarrow student(x)], [e \in E \rightarrow [protest(e), agent(e,X)]]]$$

The annotation structures defined by the QuantML abstract syntax can be deeply nested, since participation link structures contain the entity structures that they link; see the argument of the I_Q function in (11). (Their XML representations, by contrast, are ‘flat’, which is more convenient for their practical use.) A participant entity structure inside a link structure can itself have a complex structure, for instance due to the head noun of an NP being modified by a quantifying relative clause. In a well-formed annotation structure for a clause that contains only a single NP, like (7), such a link structure contains all the semantic information. The only scoping in such cases is between the NP quantifier and the verb viewed as an event quantifier (which is useful for examples like “*All passengers died in the crash*” and “*Mary wants to buy an inexpensive coat*”, cf. Szabolcsi, 2010). For clauses with multiple NPs the additional information about their relative scopes is taken into account in the I_Q function by applying ‘scoped merge’ operations to their interpretations, and where appropriate inversion operations in order to obtain the interpretations of ‘inversely linked’ quantified head noun modification by a PP or a relative clause (Barker, 2014). The reader is referred to Bunt (2018) for details.

6 Using QuantML

This section illustrates the use of QuantML with a few examples. The first example concerns two quantifications with unspecific distributivity and an NP head with adjectival modification.

- (13) a. Two young men carried all the boxes upstairs.

Markables: m1=Two young men, m2=young, m3=young men, m4=men, m5=carried upstairs, m6=all the boxes, m7=boxes

- b. QuantML annotation structure:

```
< { <m5, <carry-upstairs> }, { <m1, <<m4,man>, <m2,young>>, λz. |z| = 2, indef } }, <m6, <<m7,box0>, all, def } } },
{ <<m5, <carry-upstairs>>, { <m1, <<m4,man>, <m2,young>>, λz. |z| = 2, indef } }, agent, unspecific, narrow },
{ <m5, <carry-upstairs>>, { <m6, <<m7, box0>, all, def } }, theme, unspecific, narrow } }, { }
```

- c. Annotation representation:

```
<entity xml:id="x1" target="#m1" domain="#x2" involvement="2" definiteness="indef"/>
<qDomain xml:id="x2" target="#m3" source="#x3" restrs="#r1"/>
<sourceDomain xml:id="x3" target="#m4" individuation="count" pred="man"/>
<adjMod xml:id="r1" target="#m2" distr="individual" pred="young"/>
<event xml:id="e1" target="#m5" pred="carry-up"/>
<entity xml:id="x4" target="#m6" domain="#x5" involvement="all" definiteness="def"/>
<sourceDomain xml:id="x5" target="#m7" individuation="count" pred="box"/>
<participantLink event="#e1" participant="#x1" semRole="agent" distr="unspecific"
  eventScope="narrow"/>
<participantLink event="#e1" participant="#x1" semRole="agent" distr="unspecific"
  eventScope="narrow"/>
<scopeLink arg1="#x4" arg2="#x1" scopeRel="wider"/>
```

- d. Semantics:

$$[X, Y, E \mid e \in E \rightarrow \text{carry-up}(e), x \in X \rightarrow [\text{man}_0(x), \text{young}(x)], |X|=2, y \in Y \leftrightarrow \text{box}_0(y), x \in \text{man}^* \rightarrow [y, e \mid e \in E, y \in Y^*, \text{agent}(e,x), \text{theme}(e,y)], y \in \text{box}_0^* \rightarrow [x, e \mid e \in E, x \in X, \text{agent}(e,x), \text{theme}(e,y)]]$$

The next example illustrates the annotation and interpretation of a cumulative quantification with a mass NP. The cumulativity means that none of the two NP quantifiers outscopes the other; this is annotated as their scopes being equal.

- (14) a. The girls ate most of the chocolate.

Markables: m1=The girls, m2=girls, m3=ate, m4=most of the chocolate, m5=chocolate

- b. QuantML annotation structure:

```
< { <m3, <eat> }, { <m1, <girl, all, def> } }, { <m4, <chocolate, most, def> } },
{ <m3, <eat> }, { <m1, <girl, all, def> } }, agent, unspecific, narrow },
<m3, <eat> }, { <m4, <chocolate0, most, def> } }, theme, unspecific, narrow } }, { }
```

- c. Annotation representation:

```
<entity xml:id="x1" target="#m1" domain="#x2" involvement="all" definiteness="def"/>
<sourceDomain xml:id="x2" target="#m2" individuation="count" pred="girl"/>
<entity xml:id="x3" target="#m4" domain="#x4" involvement="most" definiteness="def"/>
<sourceDomain xml:id="x4" target="#m5" individuation="mass" pred="chocolate"/>
<event xml:id="e1" target="#m3" pred="eat"/>
<participantLink event="#e1" participant="#x1" semRole="agent" distr="individual"
  eventScope="narrow"/>
<participantLink event="#e1" participant="#x3" semRole="patient" distr="individual"
  eventScope="narrow"/>
<scopeLink arg1="#x1" arg2="#x3" scopeRel="equal"/>
```

d. Semantics:

Interpretation of the entity structure for “*most of the chocolate*”:

$$I_Q(\varepsilon_{P2}) = [X \mid |\Sigma(X)| \geq |\Sigma(\text{chocolate}_0)|/2, x \in X \rightarrow \text{chocolate}_0(x)].$$

Interpretation of the annotation structure for the sentence:

$$I_Q(\langle \varepsilon_E, \{\varepsilon_{P1}, \varepsilon_{P2}\}, \{L_{P1}, L_{P2}\} \{sc1\} \rangle) = I_Q(L\varepsilon_E, \varepsilon_{P1}) \cup I_Q(L\varepsilon_E, \varepsilon_{P2}) = \\ [X Y E \mid x \in X \leftrightarrow \text{girl}_0(x), y \in Y \rightarrow \text{chocolate}(y), |\Sigma(Y)| \geq |\Sigma(\text{chocolate}_0)|/2, e \in E \rightarrow \text{eat}(e), \\ x \in X \rightarrow [e y \mid e \in E, y \in Y, \text{agent}(e,x), \text{theme}(e,y)] \\ y \in Y \rightarrow [e x \mid e \in E, x \in X, \text{agent}(e,x), \text{theme}(e,y)]]$$

In words, this DRS says that there is a set Y of quantities of chocolate that together make up all the contextually distinguished chocolate (referred to as “*the chocolate*”), and a set of eat-events such that each of the girls in the set (X) of contextually distinguished girls (“*the girls*”) ate some of the quantities of chocolate, and each of the quantities of chocolate was eaten by one of those girls.

The third example concerns a proper name and a quantification with “count/parts” individuation and collective (!) distributivity.

(15) a. Tarzani lifted twenty-seven-and-a-half pizzas.

Markables: m1=Tarzani, m2=lifted. m3=twenty-seven-and-a-half pizzas, m4=pizzas

b. QuantML annotation structure:

$\langle \langle m2, \langle \text{lift} \rangle \rangle, \{ \langle m1, \langle \langle \text{tarzani}, \text{count} \rangle \text{ all}, \text{def}, 1 \rangle \rangle, \langle m3, \langle \langle \text{pizza}, \text{count/parts} \rangle, \langle 27.5, \text{pizza} \rangle, \text{indef} \rangle \rangle \} \rangle,$
 $\{ \langle m2, \langle \text{lift} \rangle \rangle, \{ \langle m1, \langle \langle \text{tarzani}, \text{all}, \text{def}, 1 \rangle \rangle \}, \text{agent}, \text{individual}, \text{narrow} \rangle,$
 $\langle m2, \langle \text{lift} \rangle \rangle, \{ \langle m3, \langle \langle \text{pizza}, \text{count/parts} \rangle, \langle 27.5, \text{pizza} \rangle, \text{indef} \rangle \rangle, \text{theme}, \text{collective}, \text{narrow} \rangle \}, \{ \}$

c. Annotation representation:

```
<entity xml:id="x1" target="#m1" domain="#x2" involvement="all" definiteness="de" size="1"/>
<sourceDomain xml:id="x2" target="#m1" individuation="count" pred="tarzani"/>
<event xml:id="e1" target="#m2" pred="lift"/>
<entity xml:id="x3" target="#m3" domain="#x4" involvement="27.5" definiteness="indef"/>
<sourceDomain xml:id="x4" target="#m4" individuation="count/parts" pred="pizza"/>
<participantLink event="#e1" participant="#x1" semRole="agent" distr="individual" eventScope="wide"/>
<participantLink event="#e1" participant="#x3" semRole="theme" distr="collective" eventScope="wide"/>
<scopeLink arg1="#x1" arg2="#x3" scopeRel="wider"/>
```

d. Semantics:

$$[E \mid e \in E \rightarrow \text{lift}(e), e \in E \rightarrow [Y \mid y \in Y \rightarrow \text{pizza}^+(y), |\Sigma Y|^{\text{pizza}}=27.5, \text{agent}(e, \text{tarzani}_0), \text{theme}(e, Y)]]$$

7 Conclusions and Further Work

The QuantML annotation scheme was recently proposed to the International Organisation for Standardisation for developing into part 12 of the ISO Semantic Annotation Framework, and was accepted as such in March 2019 (ISO, 2019a). The QuantML scheme is rooted in the theory of generalized quantifiers, neo-Davidsonian event semantics, and DRT, and is methodologically shaped after the ISO principles of semantic annotation (ISO standard 24617-6). Different from these semantic theories, the proposed annotation scheme has a number of provisions for leaving aspects of quantification unspecified, on the one hand intended to reflect the vagueness and ambiguity that natural language quantifications may have, and on the other hand to allow annotators to make annotations with varying degrees of granularity.

The current proposal still has several loose ends, e.g. related to modality and polarity and to intensional contexts. It is also fair to say that, where GQT and DRT do not provide adequate solutions for all the complexities of quantification in natural language, currently no annotation scheme can be expected to do much better.

The next important thing after or while further elaborating the proposed annotation scheme, is to apply it in annotation projects and see to what extent it may need to be adapted in order to be optimally useful for language technology applications and for empirically-based semantic investigations.

References

- Abzianidze, L., J. Bjerva, K. Evang, H. Haagsma, R. van Noord, P. Ludmann, D.-D. Nguyen, and J. Bos (2017). The Parallel Meaning Bank: Towards a multilingual corpus of translations annotated with compositional meaning representations. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Valencia, Spain, pp. 242–247. Association for Computational Linguistics.
- Alshawi, H. (1992). *The Core Language Engine*. Cambridge, Mass.: MIT Press.
- Barker, C. (2014). Scope. In S. Lappin and C. Fox (Eds.), *The Handbook of Contemporary Semantic Theory*, pp. 40–76. John Wiley.
- Barwise, J. and R. Cooper (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy* 4, 159–219.
- Benthem, J. v. and A. ter Meulen (1985). *Generalized Quantifiers in Natural Language*. Dordrecht: Foris.
- Bos, J. (1995). Predicate logic unplugged. In *Proceedings 10th Amsterdam Colloquium*, Amsterdam, pp. 133–142. ILLC.
- Bos, J., V. Basile, K. Evang, N. Venhuizen, and J. Bjerva (2017). The Groningen Meaning Bank. In N. Ide and J. Pustejovsky (Eds.), *Handbook of Linguistic Annotation*, pp. 463–496. Berlin: Springer.
- Bunt, H. (1985). *Mass Terms and Model-Theoretic Semantics*. Cambridge University Press.
- Bunt, H. (2017). Towards interoperable annotation of quantification. In *Proceedings ISA-13, Thirteenth International Workshop on Interoperable Semantic Annotation*, Montpellier, France.
- Bunt, H. (2018). *Semantic Annotation of Quantification in Natural Language*. TiCC Technical Report 2018-15. Tilburg Center for Cognition and Communication, Tilburg University.
- Bunt, H. and R. Muskens (1999). Computational Semantics. In H. Bunt and R. Muskens (Eds.), *Computing Meaning, Vol. 1*, pp. 1–32. Dordrecht: Kluwer Academic Publishers.
- Bunt, H., V. Petukhova, A. Malchanau, A. Fang, and K. Wijnhoven (2018). The DialogBank: Dialogues with interoperable annotations. *Language Resources and Evaluation 2018*.
- Bunt, H. and J. Pustejovsky (2010). Annotating temporal and event quantification. In *Proceedings ISA-5, Fifth International Workshop on Interoperable Semantic Annotation*, pp. 15–22. City University of Hong Kong.
- Bunt, H., J. Pustejovsky, and K. Lee (2018a). Towards an ISO Standard for the Annotation of Quantification. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan.
- Bunt, H., J. Pustejovsky, and K. Lee (2018b). Towards an ISO Standard for the Annotation of Quantification. In *Proceedings 11th International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. ELRA.
- Champollion, L. (2015). The interaction of compositional semantics and event semantics. *Linguistics and Philosophy* 38 (1), 31–66.
- Davidson, D. (1967). The Logical Form of Action Sentences. In N. Resher (Ed.), *The Logic of Decision and Action*, pp. 81–95. Pittsburgh: The University of Pittsburgh Press.

- Hobbs, J. and S. Shieber (2004). An algorithm for generating quantifier scopings. *Computational Linguistics* 13(1-2), 47–63.
- ISO (2012). *ISO 24617-1: 2012, Language Resource Management - Semantic Annotation Framework (SemAF) - Part 1: Time and events*. Geneva: International Organisation for Standardisation ISO.
- ISO (2014a). *ISO 24617-4: 2014, Language Resource Management - Semantic Annotation Framework (SemAF) - Part 4: Semantic roles*. Geneva: International Organisation for Standardisation ISO.
- ISO (2014b). *ISO 24617-7: 2014, Language Resource Management - Semantic Annotation Framework (SemAF) - Part 7: Spatial information*. Geneva: International Organisation for Standardisation ISO.
- ISO (2015). *ISO 24617-6:2015, Language Resource Management - Semantic Annotation Framework (SemAF) - Part 6: Principles of semantic annotation*. Geneva: International Organisation for Standardisation ISO.
- ISO (2016). *ISO 24617-8:2016, Language Resource Management - Semantic Annotation Framework (SemAF) - Part 8: Semantic relations in discourse, Core annotation scheme (DR-Core)*. Geneva: International Organisation for Standardisation ISO.
- Kamp, H. and U. Reyle (1993). *From Discourse to Logic*. Dordrecht: Kluwer Academic Publishers.
- Parsons, T. (1990). *Events in the Semantics of English: A Study in Subatomic Semantics*. Cambridge, MA: MIT Press.
- Pustejovsky, J., H. Bunt, and K. Lee (2010). ISO-TimeML. In *Proceedings of LREC 2010*, Malta. ELDA, Paris.
- Pustejovsky, J., H. Bunt, and A. Zaenen (2017). Designing annotation schemes: From theory to model. In N. Ide and J. Pustejovsky (Eds.), *Handbook of Linguistic Annotation*, pp. 21–72. Berlin: Springer.
- Quine, W. (1960). *Word and Object*. Cambridge, Mass.: MIT Press.
- Schwertel, U. (2005). *Plural Semantics for Natural Language Understanding - A Computational Proof-Theoretic Approach*. Ph.D. Thesis, University of Zurich.
- Szabolcsi, A. (2010). *Quantification*. Cambridge (UK): Cambridge University Press.
- Westerståhl, D. (1985). Determiners and context sets. In J. van Benthem and A. ter Meulen (Eds.), *Generalized Quantifiers in Natural Language*, pp. 45–71. Dordrecht: Foris.
- Willis, A. and S. Manandhar (2001). The availability of partial scopings in an underspecified semantic representation. In *Computing Meaning, Vol. 2*, pp. 129–145. Dordrecht: Springer.