

Representing Verbal Semantics with Diagrams

An Adaptation of the UML for Lexical Semantics

Andrea C. SCHALLEY

School of Languages, Cultures and Linguistics
University of New England
Armidale, NSW 2351, Australia
andrea.schalley@une.edu.au

Abstract

The paper presents a new way of accounting for the meaning of verbs in natural languages, using a diagrammatic notation based on the Unified Modeling Language (UML). We will introduce the new framework by outlining some modeling elements and indicating major differences to the UML. An extended example will be discussed in more detail. We will then focus on the cognitive background of the framework, and in particular address the question why the usage of graphical elements within a linguistic modeling language proves to be very fruitful. Finally, we will briefly indicate the potential of the new framework and its applicability.

1 Introduction

Today, the Unified Modeling Language (UML) is accepted as lingua franca for the design of object-oriented systems, being widely used for software development processes. Although the UML has also been employed in other fields such as business modeling (cf. the example profile in the UML specification, Object Management Group 2003), research in theoretical and computational linguistics has not yet tried to apply a graphical language as rich as the UML. Such an approach will be advocated in this paper, focussing on the question how verbal meaning is to be represented adequately. Our answer is a new framework adapted from the UML to model the meaning of verbs, as developed in extenso in Schalley (2004). This framework for linguistic semantics is called *Unified Eventivity Representation* (UER), because it is a true extension of the UML and not just a profile. Living up to its name, the UER tries to unify both intuitivity and formality. It employs intuitive semantic primes as building blocks and includes these within a framework of specified modeling elements. Semantics and syntax of the modeling elements are explicated in the specification of the UER, which was done in the style of the UML specification. Being an adaptation from the UML, the UER introduces a third formal paradigm of computer science into linguistic semantics, one that is

neither functional nor logical but object-oriented in nature. This is one of the factors contributing to the cognitive adequacy of the UER. Since the UER is based on the UML, it can be easily put to use in computational linguistics.

In Section 2, the UER as adapted from the UML will be sketched. The cognitive relevance of the UER is outlined in Section 3 — graphical modeling elements in general represent prominent kinds of concepts, or, respectively, structural or meta-concepts. In particular, we will discuss the importance of these graphical modeling elements for the cognitive adequacy of representational frameworks such as the UER. Section 4 finally comprises an outlook, listing some potential areas of application of the UER.

2 A Diagrammatic Modeling Language for Linguistic Semantics

The UER exhibits a novel use of a diagrammatic notation. But even though it adapts a well-known and well-elaborated framework used in computer science, it constitutes not only a new use of the UML's diagrammatic notation but is a redesign in order to achieve better cognitive and in particular linguistic adequacy. In order to get a grasp of the character of the UER we will indicate some major differences between the UER and the UML and then discuss an extended example, modeling two particular concrete readings of a verb in more detail.

The UER's focus does not rest on computational adequacy to the same extent as the UML's does. It is not designed to develop software systems, but to represent meaning. This objective entails, for example, that pure software specific modeling elements are not part of the UER. Instead, it aims to be close to conceptualization as revealed in natural language. Nevertheless, the UML has been a perfect starting point for the endeavour of developing an adequate modeling language for verbal semantics, because on a coarse level of granularity the UML itself supports cognitive modeling in the sense that it allows

to model software requirements, without going too deep into implementational issues in the beginning. But why has a new framework been developed — why has it not been sensible or feasible to establish a UML profile? We believe that for the purpose of a linguistic modeling language which is designed to represent verbal semantics, new modeling elements are required and therefore new metamodel elements are inevitable.

Roughly, verbs encode events or similar entities, entities that are called eventities within the UER.¹ That is, the semantics of a verb or, as in most cases, one of its readings, corresponds to the eventity that is encoded by the verb.² Then, to represent the meaning is to model the eventity.

We believe that eventities are conceptual units in human cognition and comprise particular components that are combined in a particular way. In order to model this appropriately, the UER has a graphical container for eventity concepts: octagons represent eventities as such and contain their components in a structured way. As modeling elements of the UER, they are called eventity frames and model eventities, the specificities of which are graphically contained within these diagrammatic elements. Similarly, the components are again diagrammatic elements as long as they represent structural or meta-concepts that hold content. Here the notions of state and transition (which are part of the UML), or the new notion of participant class, which is a modeling element reminiscent of the UML's classifier role, could be mentioned. New modeling elements not being part of the UML, such as the eventity frame or participant class, clearly establish an extension of the UML. Therefore, the UER is not a profile of the UML, but a close relative.

In addition to defining new modeling elements for the UER (and adapting UML ones), the UML's division into different modeling views resulting in separate diagrams (such as class, statechart, activity, or collaboration diagrams, cf. Object Management Group 2003: I-2)³ is given up in the UER. Since both dynamic as well as static aspects are part

of an eventity concept (in modeling an eventity one has to answer the questions what is happening and to whom it is happening), both aspects are modeled within one eventity frame and not in several separate diagrams. Nevertheless, we take care to distinguish the aspects within the eventity frame, with the dynamic aspects being graphically contained by the so-called dynamic core, thereby forcing users of the UER to go for clear distinctions. The integration of dynamic and static aspects seems to be feasible in the UER, because — differently from software systems — we expect eventity concepts not to exceed a particular level of complexity due to an assumed maximum of complexity applying to any cognitive unit. Accordingly, we expect eventity frames to only come up with straightforward models which are in general easily manageable.

To get a better idea what a model of an eventity, or, respectively, a model of a verb's reading looks like, we will discuss two readings of *wake up* in the following. Consider the eventity frame representing the semantics of its non-causative reading in Fig. 1 (as in *He woke up* or in *Suddenly Eleni woke up*).

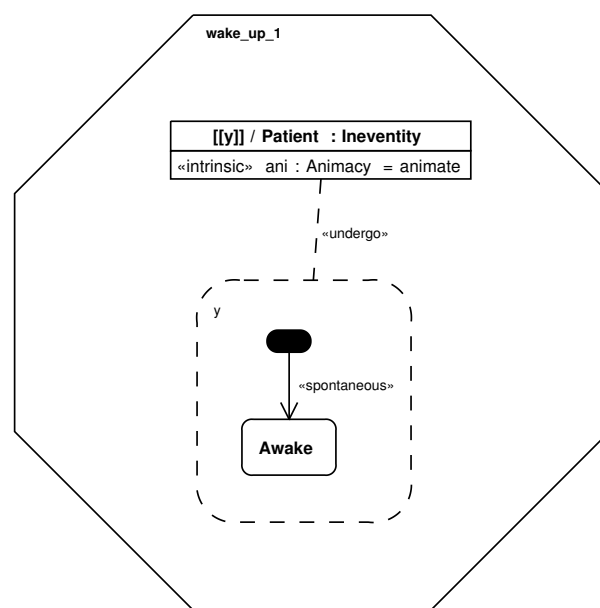


Figure 1: The non-causative reading of *wake up*

The octagon depicts the eventity that is encoded by the verb and hence represents the conceptual unit corresponding to the verb's reading. Each eventity frame can have a name shown in the upper left corner, in this case the name *wake_up_1* was selected according to the first 'wake up' eventity, the WAKE_UP_1 eventity. The eventity's components are nested within the octagon. First of all, participants of an eventity are rendered as solid-outline rectangles similar to UML class symbols and are

¹ The term 'eventity' has been adopted from Zaefferer (2002).

² We use the term eventity as a term for a kind or type, not an instantiation or token. Hence, 'to wake up' is an eventity, termed WAKE_UP_1 in the following (eventities are usually notated with capital letters), whereas 'John wakes up' is an instantiation of this eventity.

³ Not all parts of the UML are relevant to our endeavour. Accordingly, non-relevant parts such as use case and implementation diagrams are not included into the specification of the UER at all in order not to overload it and to adjust it to our purposes.

attached to the dynamic core by a dashed line (as in UML collaboration diagrams) that indicates their status as participants, the participate association. In the case of WAKE_UP_1, there is only one participant, the undergoer that wakes up and thus endures the course of events.⁴ The participant's name expression, $[[\underline{y}]]$, means that \underline{y} is a representative name for entities that could potentially be participants in the eventity (therefore, a notation reminiscent of mathematical equivalence class notation has been chosen).

The dynamic components of the eventity concept are contained in the dynamic core itself which is displayed as a dashed-outlined rectangle with rounded corners and generally comprises the state machines of prominent participants.⁵ For the sake of clarity (and because it is necessary when there is more than just one participant), the representative's name, in this case \underline{y} , is cited in the upper left corner of its state machine. \underline{y} experiences a transition from an unknown source state to the known passive simple state of being *Awake*. In order to undergo a transition into the state of being awake and thus to undergo a change of state at all, the undergoer must have been in a source state differing from the target state. This is the only information we do have, we do not know whether the undergoer was asleep, dozing or unconscious while in the source state. But \underline{y} must have been in a state which was not the one of being awake, such that a transition could result in the state of being awake. Accordingly, the target state is specified, whereas the source state is displayed as an unspecified source state. Moreover, there is no reason and thus trigger for the transition conceptualized, therefore the transition is marked as \ll spontaneous \gg and distinguished from completion transitions or transitions triggered by signals.

\underline{y} in the dynamic core is a reference to the participant class displayed outside of the dynamic core, in the static periphery. The static periphery in general depicts the participants, the roles they hold within the eventity, and the relationships that hold between them. In the example modeling in Fig. 1, the undergoer \underline{y} is in fact a patient, i.e., has the role of undergoing some change of state or condition (and not only a change of location as a theme, for in-

⁴ The term *undergoer* — as well as its counterpart *actor* — are taken from Van Valin and LaPolla (1997: 141–147).

⁵ There are at most two prominent participants in each eventity, the most active one (the *actor*, with the stereotype \ll do \gg attached to its participate association) and the most passive one (the *undergoer*, with the stereotype \ll undergo \gg attached to its participate association). Only prominent participants are assigned their own state machine within the dynamic core.

stance, would). This is indicated in the role specification *Patient* within \underline{y} 's participant class. Additionally, in order to be 'wakeable' (that is to say, to be a potential undergoer of a WAKE_UP_1 eventity), \underline{y} has to be an *animate Ineventity*, a non-eventity entity in the UER's participant ontology which has the intrinsic property of being animate. The specification of both an ontological category as well as a property (technically captured in the UER in form of a type-expression and an attribute) rules out that SLEEP, as an eventity, or a stone, as an inanimate ineventity, could be potential undergoers of the WAKE_UP_1 eventity. In other words, the participant class specification works like a filter on the set of all entities, ruling out those entities which cannot be participants of the eventity in question because they do not fulfil the required characteristics. This is important in describing verbal semantics because it is an essential part of a verb's semantics which selectional restrictions apply.

The first, non-causative reading of *wake up* is simple in the sense that only one participant is involved. Turning to the causative reading, WAKE_UP_2 (as in *He woke me up* or *The storm woke him up*), the modeling in Fig. 2 becomes more complex: there are two interacting participants, both of which are prominent participants — they are both assigned their own state machine, each of which is depicted in a swimlane in the dynamic core (the two swimlanes are divided by a solid vertical swimlane border).

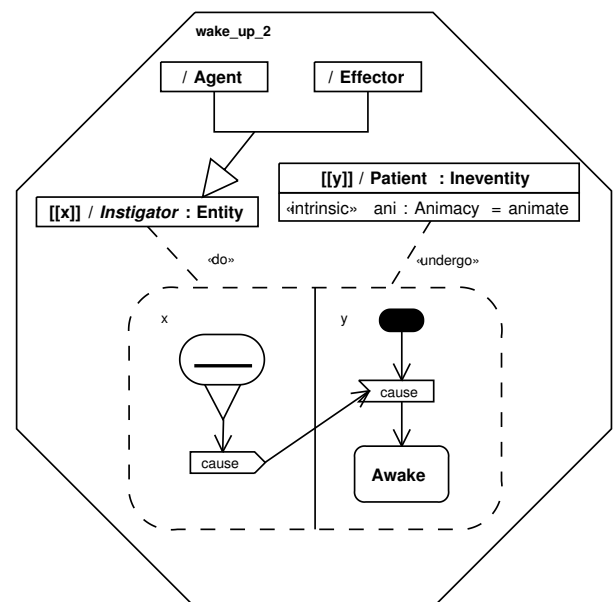


Figure 2: The causative reading of *wake up*

The undergoer \underline{y} essentially endures the same course of events, although this time the transition

has a trigger in that it is caused by the active instigator, the actor. That is, the transition is triggered by a cause signal, the receipt of which is rendered in the concave pentagon in y 's swimlane. The signal is sent from the actor x , with the signal sending being represented in x 's swimlane as convex pentagon. The signal sending is the result of an unspecified action state. In other words, the actor performs some action (where action is broadly understood and does not necessarily entail intention), the nature of which is not conceptualized and irrelevant, thus leaving the specifics of the action state irrelevant (which is indicated by the underscore). All that is important is that there is some 'action' by x involved so that x wakes y up.⁶

In the case of the actor's specification in the static periphery, there are not many restrictions. The actor is primarily an entity that instigates the transition. The italicized role description pays tribute to the fact that natural languages distinguish between volitional and involitional instigators. *Instigator* is an abstract role description, meaning that it cannot be directly instantiated but only by its children (we employ the object-orientational concept of inheritance at this point), namely *Agent* (volitional instigator) or *Effector* (involitional instigator). Although in English users are not forced to decide whether they are conceptualizing an agent or an effector, in Hare, for example, an Athapascan language spoken in Canada, users have to mark agents differently than effectors (cf. Frawley 1992: 207). Thus, Hare forces its users to definitely make a decision whether it is an agent or an effector they are talking about. If the model was a modeling of Hare (and not English as it is), we would add the constraint $\{\text{disjoint}\}$ to the inheritance relations in order to indicate that the actor can be either an agent or an effector, but not both at the same time. This way, natural language specificities come into the models.

3 Cognitive Relevance of Diagrammatic Modeling Elements

We trust that the above illustrations suffice to give an impression what UER diagrams entail, although there are details of the diagrams that have not been explained. But in exactly what way does such a graphical representation as in Fig. 2 differ from a Wunderlich-style decomposition as in (1) (cf. also Wunderlich 1996, 1997)?

(1) CAUSE(x , BECOME(AWAKE(y)))(s)

⁶ Action states are rendered with convex arcs on both sides, whereas passive states are shown as rectangles with rounded corners.

Essentially, the same information concerning dynamic structuring seems to be included in (1): there is an x which causes that y becomes awake. But obviously necessary information about the participants is not included. Of course, one could add this information as in (2) (in this case, the decomposition is not within Wunderlich's framework any more).

(2) (AGENT(x) \vee EFFECTOR(x)) \wedge
 PATIENT(y) \wedge ANIMATE(y) \wedge
 INEVENTITY(x) \wedge INEVENTITY(y) \wedge
 CAUSE(x , BECOME(AWAKE(y)))(s)

Comparing (2) to Fig. 2, the diagrammatic representation is to be preferred for several reasons, one of these being the intuitivity that is brought forward in the graphical structure: those modeling elements that are cognitively connected are graphically connected via connectors, containment resp. nesting, or visual attachments. That is, the cognitive structuring is reflected in the diagrammatic representation in a straightforward way, which is not the case in the linearized representation in (2). Moreover, the explicit partition of static and dynamic aspects within one eventivity frame as well as the specified syntax of the modeling elements facilitates not only the understanding of the representation, but at the same time forces users of the UER to make sure they produce sound diagrams. That entails reflecting on what exactly causes, for instance, the undergoer's transition in WAKE_UP_2. In (1) and (2) it is x as a participant that directly causes the transition, while in Fig. 2 it is some action of x that causes the transition. The latter is more appropriate and also supported by speakers' intuition and conceptualization — something has, in a very broad sense, to 'happen' (even if it was pure presence) in order to cause a transition.⁷ To represent that 'something', a feature has been included into the UER that is not part of the UML, namely unspecified elements (generally rendered with underscores in the name slot). These are elements where only the graphical — and thus cognitive — structure is present (such as the action state in the actor's swimlane in Fig. 2), but no content of the structure is given. That is, the exact concept is irrelevant and underspecified, the only thing that matters is structure: in Fig. 2 it is merely conceptualized that some kind of action takes place, but

⁷ This is also supported by the fact that (a) is fine, but (b) is not:

(a) *The ball broke the window.*

(b) **The window broke as a result of the ball.*

In other words, explication (b) would need some action of the ball to be specified in order to be sound, such as in (c):

(c) *The window broke as a result of the ball's rolling into it.*

the specifics of the action are not specified. Since the UER aims at representing cognitive structures, this is a sensible feature — which it would not be within the UML, of course, as the UML has to head towards determinism, being a computational modeling language.

The general layout of both the UML and the UER as graphical languages supports cognitive adequacy. In graphical languages such as the UML and the UER, prominent structural concepts can be distinguished by non-textual symbols, namely their corresponding graphical modeling elements. Examples in the UER we have seen are the octagon for eventities, the rectangle for participants, and the rectangle with rounded corners for passive simple states. In other words, the structure is passed on into the graphical-visual domain, whereas the contents are kept in linearized form. Representational languages in linguistic semantics to date only rarely distinguish in their formalizations between structural meta-concepts, such as the ones mentioned above, e.g., or state and transition, and the concepts themselves, such as the state of being *Awake*. Meta-concepts exist at the most implicitly in the arity of predicates, in particular predicate names (where the reader has to know that *BECOME*, for example, represents a transitional concept), or in fixed combinations of predicates. Although meta-concepts are extensively discussed in the literature (cf. the discussions about *Aktionsarten*), it is not taken care to explicate them in a distinct way and to distinguish them from specified concepts, i.e. their contents.

The UER is the first linguistic representational framework that explicitly accounts for meta-concepts, rendering them graphically and thus fundamentally different from their contents, and displaying different structural concepts with visually different modeling elements. That way, an intuitive line is drawn that divides these two levels of representation, that divides the specified concepts embedded in a structure from the structure itself. The distinction between structural concepts and ‘filled’ concepts is in our eyes a very vital one not only in terms of modeling precision — users have to clearly distinguish these levels of representation — but also in terms of cognitive adequacy and universality. We believe that structural meta-concepts are universal due to them being based on human experiences, and that they are in principle stable. This does most certainly not apply in this generality to contents, although one might assume a very fundamental shared knowledge across cultures and thus languages.

Hence, in fixing the meta-concepts but not their contents, the UER is a modeling language that can readily accommodate different linguistic facts and allows for adequate recording of language specificity due to its flexibility concerning the ‘filling’ of the meta-concept. Since it moreover includes the UML’s extension mechanisms such as constraints and stereotypes, it can even be more easily adapted to modeling needs — modeling elements and thus cognitive components can, if necessary, be adjusted. The degree of granularity of a representation and thus the understanding of what is primitive within a modeling can be aligned to the modeling purposes. It is our hope that the UER is a modeling language that can be universally applied to model verbal semantics because of its general flexibility, no matter what natural language’s verbal system is described and what granularity is needed for the semantic descriptions.

4 Applicability in Linguistic Research

This last section is devoted to an outlook on what the UER should be able to do and where potential applications of this new modeling language can possibly be found. Within linguistic semantics, the UER is expected to facilitate and enhance research. First of all, as has already been indicated, it forces semanticists to reflect on their representations, in particular on the structure they model and on what they consider to be primitive components within their endeavour. The syntax of the UER has been specified with the aim to allow sensible linguistic modeling; ensuring that diagrams are syntactically correct (that is, conform to the specification) will be of invaluable help in semantic work, and the design of the UER allows, as has been argued, to achieve cognitive adequacy more easily than other rigorous linguistic frameworks do. Secondly, the UER supplies mechanisms to capture any potential component of verbal semantics, thus allowing for comprehensive modeling. Thirdly, case studies applying the UER framework have shown a strong potential of the UER in capturing structural relations between different readings of one lexical item and also between different lexical items. In other words, the UER is a practical tool for the study of semantic relationships.

In the study of polysemy, that is, the study of lexical items (elements of our lexicon) and their different readings shared on systematic grounds, explicating the different readings in the UER allows to pin down the systematic changes from one reading to another (cf., for instance, Chapter 9 in Schalley 2004). A comparison of UER diagrams and intu-

itions of native speakers in cases of polysemy has shown a thrilling interconnection. The closer the modeling of the different readings were structurally to each other, the surer were native speakers intuitively that the readings in questions were instances of polysemy. On the other hand, it seems as if at least a major change affecting one of the graphical modeling elements has to occur — such as the gain of a participant (cf. the two readings of *wake up* as modeled in Fig. 1 and 2) — in order for native speakers to readily identify different readings of a lexical item and not to consider instances of both readings to be instantiations of only one reading used in different contexts.

Taking the distinction between structure and content into account, the UER offers a new perspective on decompositional semantics. In eliminating all content within modeled readings and just keeping the graphical structure, one can ask what changes the remaining structures can undergo and whether one finds instances of such changes in the semantics of natural languages. This enables scholars to systematically study not only polysemy, but also semantic change and verb classification, and to determine where irregularities in meaning structures of natural languages are to be found. Moreover, alternations such as causativation or resultatives can be systematically captured and studied within the UER.

In addition, it might be interesting to ask what degree of complexity eventivity structures can maximally obtain while still being lexicalized in a single verb. In other words: what are possible verbs? This can be investigated within the UER framework, because the UER in principle allows for the modeling of eventivities which are too complex to be coded in a single verb. For instance, one would expect that, if there are two prominent participants involved (i.e., actor and undergoer), there has to be some interaction between the participants, some causation to take place in order for the modeled eventivity to be lexicalized. A verb expressing that something becomes liquid and another something becomes full at the same time is not likely to exist (also cf. Kaufmann 1995: 198f.). Systematic analysis applying the UER can show which structures are most likely ruled out because they do not constitute a cognitive unit, with the missing unity showing up in unconnected swimlanes within the dynamic core, for instance.

There are, due to the proximity of the UER to the UML, not only theoretical linguistic, but naturally computational linguistic applications that sug-

gest themselves. For example, the usefulness of UER structures for machine translation could be explored. Pattern matching could be applied to UER diagrams. Having the semantics of a verb in the source language captured in a diagram, one could ask which verb and thus which representation diagram in the target language would be most appropriate for the translation.⁸ This should be the verb the representation of which comes ‘closest’ to the modeled one in the source language. Evidently, criteria for what is considered to be ‘closest’ would have to be identified.

Apart from progress in scientific discovery, the UER as a UML-based modeling language is a new modeling language that allows for envisaging applications also in speech processing systems, for example. The UER is a rigorous, but cognitively oriented, non-iconic, but intuitive decompositional modeling language both suitable for human and computational usage. We hope that the UER will be tested extensively and applied to different research areas within theoretical and computational linguistics. Since it is a rather new modeling language, its testing is just about to begin, but we are confident that the UER will prove very fruitful for many research enterprises.

References

- William Frawley 1992. *Linguistic Semantics*. Hillsdale: Lawrence Erlbaum.
- Ingrid Kaufmann 1995. *Konzeptuelle Grundlagen semantischer Dekompositionsstrukturen. Die Kombinatorik lokaler Verben und prädikativer Komplemente*. (Linguistische Arbeiten 335.) Tübingen: Niemeyer.
- Object Management Group 2003. *OMG Unified Modeling Language Specification*. Version 1.5.
- Andrea C. Schalley 2004. *Cognitive Modeling and Verbal Semantics. A Representational Framework Based on UML*. (Trends in Linguistics. Studies and Monographs 154.) Berlin/New York: Mouton de Gruyter.
- Robert D. Van Valin and Randy J. LaPolla 1997. *Syntax: Structure, Meaning and Function*. (Cambridge Textbooks in Linguistics.) Cambridge: Cambridge University Press.
- Dieter Wunderlich 1996. Models of lexical decomposition. *Lexical Structures and Language Use. Proceedings of the International Conference on Lexicology and Lexical Semantics*, ed.

⁸ This is of course a simplification of the translation problem, because idiomatic constructions and syntactic environments are, for instance, not taken into account — or one would have to model them within the UER as well.

Edda Weigand and Franz Hundsnurscher. Tübingen: Niemeyer, 169–183.

Dieter Wunderlich 1997. Cause and the structure of verbs. *Linguistic Inquiry* 28(1):27–68.

Dietmar Zaefferer 2002. Polysemy, polyvalency, and linking mismatches. The concept of *rain* and its codings in English, German, Spanish, and Italian. *DELTA — Documentação de Estudos em Lingüística Teórica e Aplicada* 18(spe.):27–56.