Integrating Generative Lexicon and Lexical Semantic Resources

James Pustejovsky Elisabetta Ježek

May 23, 2016

Tutorial Schedule

1. **Introduction to GL** (9:00 am - 10:30 am)

   (a) Basic concepts in Generative Lexicon

   (b) Notation and Language: typed feature structures

   (c) Qualia Structure

   (d) Events and their participants

   (e) Meaning Composition in GL
   - encoding selection
   - encoding coercion
   - encoding subselection
   - encoding cocomposition

   Readings: Pustejovsky and Jezek (2008): Coercion
   Pustejovsky and Jezek: Chapter 2, Intro to GL
   Pustejovsky Co-Compositionality

   Coffee Break (10:30 - 11:00)

2. **Enriching Lexical Resources with GL** (11:00 - 1:00 pm)

   (a) Case Study 1: Enriching VerbNet with Dynamic Event Structure

   (b) Case Study 2: Enriching AMR with Dynamic Argument Structure

   (c) Case Study 3: Enhancing WordNet Verb and Noun Ontology with Telic Relations
   Readings: Pustejovsky, Palmer, Zaenen, and Brown (2016): Verb Meaning in Context
   Fellbaum (2013): Purpose Verbs
   Banarsecu et al: Abstract Meaning Representation

Tutorial Description

In this tutorial, we demonstrate how elements of Generative Lexicon Theory (GL) can be used to help enrich both established and developing lexical and computational semantic resources within the CL community. This includes lexicons, ontologies, annotation schemes, and annotated corpora (VerbNet, AMR, and WordNet).

The tutorial is organized into two parts. The first part aims to acquaint the audience – computational linguists, natural language engineers, and language resource developers – with the basic assumptions and components of the theory and motivate theoretical decisions through evidence-based analysis over large linguistic datasets. We show how the theory models the interaction between lexical information and other components of grammar; in particular, how it mediates various problems in the mapping from lexical semantic representations to syntactic forms and, to a lesser extent, to pragmatic interpretation. We discuss the significant developments of the theory since the original statement in Pustejovsky (1995), including the elaboration of a general theory of semantic selection and semantic typing (Asher and Pustejovsky 2006, Pustejovsky 2011). Finally, we illustrate how the theory has drawn increasingly on the findings of corpus linguistics and distributional semantic analysis (Pustejovsky and Jezek, 2008, Pustejovsky and Rumshisky, 2008, Jezek and Quochi, 2010, Jezek and Vieu, 2014), creating a new dimension of evidence-based analysis and interpretation.

Some of the most difficult problems recently addressed by GL include: how to encode the dynamic interpretation of events and their participants (Pustejovsky 2013, Jezek and Pustejovsky 2015); the extension of the Telic qualia role to verbs, e.g., rationale and purpose clauses; how to distributionally model the range and effect of coercion phenomena, incorporating a CPA methodology (Hanks and Pustejovsky 2005, Hanks 2013, Jezek et al. 2014) and a broader notion of context.

In part two, we explore how the semantic phenomena illustrated in part one are implemented and handled in existing resources, examining several case studies: VerbNet, WordNet, and AMR. We demonstrate how both the representational facilities and the compositional mechanisms native to GL can simplify and extend the theoretical infrastructure of these resources. In particular, we propose enhancements to VerbNet (Palmer 2009) and AMR (Banarescu et al. 2014) leveraging the work on dynamic event structure and argument encoding presented in part 1. We then show, following a proposal in Fellbaum (2013), how WordNet verb links can be enriched with Telic qualia, to encode the purposes and goals associated with verbs.

Motivation and Topics of Interest

Recently, techniques and strategies for the acquisition of lexical semantic information for natural language resources have changed dramatically, influenced by the availability of ever-larger corpora, distributional methods, and newly annotated or semi-annotated corpora. In spite of these developments, however, researchers interested in creating lexical resources still face the problem of anchoring the selection of linguistic features used in the acquisition of information to a model which is theoretically well-developed, while overcoming common problems such as data sparsity and lexical ambiguity. Semantic feature do no always emerge from a purely corpus-based distributional analysis (Pustejovsky and Jezek 2008); moreover, there is often no consensus on what features to use for general acquisition tasks, and in many cases, the feature sets are constructed ad-hoc to address the goals of the specific task. Because GL has long approached these problems of polysemy, type coercion, metonymy, and co-composition from a systematic and theoretical perspective, it is worth examining how the theory can contribute to enriching and extending existing lexical resources which have emerged within the CL community.

GL has already been exploited as a theoretical background in language resources. Perhaps the most significant contribution of GL to computational lexicography took place in the framework of the EU-sponsored SIMPLE project (Semantic Information for Multipurpose Plurilingual Lexicons), whose aim was to develop comprehensive semantic lexicons for 12 European languages. In this context, an extended version of the Qualia Structure was proposed (Lenci et al 2000). Further, qualia structure was proposed as an organizing principle for the top ontology in EuroWordNet (Vossen 2001). GL Semantic typing has also been extensively used in the construction of PDEV (Pattern Dictionary of English Verbs, Hanks and Pustejovsky 2005), where semantic distinctions among the different senses of verbs depend on the semantic type of the arguments, as predicted by the co-composition principle, as well as in the design of the Brandeis Semantic Ontology (Pustejovsky et al 2006, Havasi et al, 2009). Finally, GL's event structure was developed into a subeventual lexical resource in Im (2013) that explores the principles of opposition structure and change in GL.

In this tutorial we make use of this background and of recent work to propose enhancements to existing resources widely used in the community. For all these reasons a tutorial illustrating how GL principles can be put into practice in linguistic analysis and lexical resource building.

Instructors

James Pustejovsky holds the TJX Feldberg Chair in Computer Science at Brandeis University, where he directs the Lab for Linguistics and Computation, and chairs both the Program in Language and Linguistics and the Computational Linguistics Graduate Program. He has conducted research in computational linguistics, AI, lexical semantics, temporal reasoning, and corpus linguistics and language annotation. He has written several books on computational semantics, computational linguistics, and corpus processing. He has authored numerous books, including *Generative Lexicon*, MIT, 1995; *Semantics and the Lexicon*, Springer, 1993; *The Problem of Polysemy*, CUP, 1996 (with B. Boguraev); *The Language of Time*, OUP, 2005 (with I. Mani and R. Gaizauskas), *Interpreting Motion: Grounded Representations for Spatial Language*, OUP, 2012 (with I. Mani), and *Natural Language Annotation for Machine Learning*, O'Reilly, 2012 (with A. Stubbs). Recent books include: *Recent Advances in Generative Lexicon Theory*, Springer, 2013; *A Guide to Generative Lexicon Theory*, OUP, Forthcoming (with Elisabetta Jezek).

Elisabetta Jezek is an Associate Professor at the University of Pavia, where she has taught Syntax and Semantics and Applied Linguistics since 2001. Her research interests and areas of expertise include lexical semantics, verb classification, theory of argument structure, event structure in syntax and semantics, lexicon/ontology interplay, word class systems, and computational lexicography. She has edited a number of major works in lexicography and published contributions focusing on the interplay between corpus analysis, research methodology, and linguistic theory. Her publications include: *Classi di Verbi tra Semantica e Sintassi*, ETS, 2003; *Lessico: Classi di Parole, Strutture, Combinazioni*, Il Mulino, 2005 (2nd ed. 2011); *The Lexicon: An Introduction*, OUP, 2015; and *A Guide to Generative Lexicon Theory*, OUP, Forthcoming (with James Pustejovsky).

References

Banarescu, Laura, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2014. *Abstract Meaning Representation (AMR) 1.2 Specification*, url = https://github.com/amrisi/amr-guidelines/blob/b0fd2d6321 ed4c9e9fa202b307cceeae36b8c25b/amr.md

Asher, Nicholas. 2011. *A Web of Words: Lexical Meaning in Context*, Cambridge University Press, Cambridge.

Asher, Nicholas and James Pustejovsky. 2006. "A Type Composition Logic for Generative Lexicon". In *Journal of Cognitive Science* 6:1-38.

Fellbaum, Christiane. 2013. "Purpose Verbs". In Pustejovsky J. et al. *Advances in Generative Lexicon Theory*, Dordrecht, Springer, 371-384.

Hanks, Patrick and James Pustejovsky. 2005. "A Pattern Dictionary for Natural Language Processing". *Revue Franaise de linguistique applique*, 10.2: 63-82.

Hanks, Patrick. 2013. Lexical Analysis: Norms and Exploitations. Cambridge Mass. The MIT Press.

Havasi, Catherine, Robert Speer, James Pustejovsky, and Henry Lieberman. "Digital intuition: Applying common sense using dimensionality reduction." *Intelligent Systems*, IEEE 24, no. 4 (2009): 24-35.

Im, Seohyun. 2013. "The generator of the event structure lexicon (GESL): automatic annotation of event structure for textual inference tasks". Ph.D. Dissertation, Brandeis University.

Jezek,Elisabetta. 2016. *The Lexicon: An Introduction*. Oxford, Oxford University Press.

Jezek,Elisabetta and James Pustejovsky. 2015. "Dynamic Argument Structure", Università di Pavia and Brandeis University, manuscript.

Jezek, Elisabetta and Magnini, Bernardo and Feltracco, Anna and Bianchini, Alessia and Popescu, Octavian. 2014. "T-PAS: A resource of corpus-derived Types Predicate-Argument Structures for linguistic analysis and semantic processing". In Calzolari N. et al. *Proceedings of the Ninth International Conference on Language Resources and Evaluation* (LREC'14), May 26-31, 2014, Reykjavik, Iceland, ELRA.

Jezek, Elisabetta and Valeria Quochi. 2010. "Capturing Coercions in Texts: a First Annotation Exercise". In Calzolari N. et al. *Proceedings of the Seventh*

International Conference on Language Resources and Evaluation (LREC'10). Valletta, Malta, May 19-21, 2010, 1464-1471, ELRA.

Jezek, Elisabetta and Laure Vieu 2014. "Distributional analysis of copredication: Towards distinguishing systematic polysemy from coercion". In Basili R., Lenci A., Magnini B. (eds.) *First Italian Conference on Computational Linguistics CLiC-it 2014* (Dec. 9-10, 2014). Pisa: Pisa University Press, 219-223.

Lenci, Alessandro and Bel, Nuria and Busa, Federica and Calzolari, Nicoletta and Gola, Elisabetta and Monachini, Monica and Ogonowski, Antoine and Peters, Ivonne and Peters, Wim and Ruimy, Nilda. 2000. "SIMPLE: A general framework for the development of multilingual lexicons." *International Journal of Lexicography* 13.4: 249-263.

Palmer, Martha. "Semlink: Linking propbank, verbnet and framenet." In *Proceedings of the Generative Lexicon Conference*, pp. 9-15. 2009.

Pustejovsky, James. 1995. *The Generative Lexicon*, MIT Press, Cambridge, MA.

Pustejovsky, James. 2011. "Coercion in a general theory of argument selection", *Linguistics*, vol. 49, no 6. de Gruyter.

Pustejovsky, James 2013. "Dymanic Event Structure and Habitat Theory", *Proceedings of GL2013*, 1-20.

Pustejovsky, James, Catherine Havasi, Jessica, Littman, Anna Rumshisky and Marc Verhagen. 2006. "Towards a generative lexical resource: The Brandeis Semantic Ontology". In Proceedings of the Fifth LREC Conference (Vol. 7).

Pustejovsky, James and Elisabetta Jezek. 2008. "Semantic Coercion in Language: Beyond Distributional Analysis", *Italian Journal of Linguistics*, 20:1, 181-214.

Pustejovsky, James and Elisabetta Jezek. Forthcoming. *A Guide to Generative Lexicon Theory*, Oxford, Oxford University Press.

Pustejovsky, James and Anna Rumshisky. 2008. "Between chaos and structure: Interpreting lexical data through a theoretical lens". In *International Journal of Lexicography* 21:3, 337-355.

Pustejovsky, James, Anna Rumshisky, Olga Batiukova, and Jessica L. Moszkowicz. 2014. "Annotation of compositional operations with GLML." In *Computing Meaning*, Dordrecht, Springer, 217-234.

Vossen, Piek. ”Tuning Document-Based Hierarchies with Generative Principles”. 2000. In Proceedings of The First International Workshop on Generative Approaches to the Lexicon (GL2001).

Contents

Chapter 1

Qualia Structure. *James Pustejovsky and Elisabetta Jezek*

To Appear from Oxford University Press
2016

A Guide to Generative Lexicon Theory

James Pustejovsky and Elisabetta Jezek

Chapter 2

Introducing Qualia Structure

In this chapter we introduce the mechanism used in GL to represent the core meaning of words. This system is based on four dimensions of meaning called *Qualia*. Qualia capture different properties of objects, as they are reflected in language: the *Formal* encodes taxonomic information, the *Constitutive* encodes information about the material and parts of objects, the *Telic* encodes information regarding the function and purpose, and the *Agentive* encodes information about the origin of an object. In the introduction of Qualia, we focus on nouns and touch only briefly on verbs in section 2.5. A thorough illustration of how Qualia can be used to represent the meaning of verbs and adjectives is developed later in Ch. 4. The goal of this chapter is to become acquainted with the Qualia formalism, to understand its motivating principles, and to be able to apply it to the analysis of novel words and expressions. Qualia Theory is introduced informally in this chapter and will be developed gradually throughout the book.

2.1 What is a Quale?

A Quale (singular of *Qualia*), from the Latin meaning "of what kind of thing", is a term GL borrows from philosophy to indicate a single aspect of a word's meaning, defined on the basis of the relation between the concept expressed by the word and another concept that the word evokes. Among the conceptual relations that a word may activate (for example, for the noun *dog*, having fur, barking, tail wagging, licking, etc.), Qualia relations as defined in GL are those that are relevant for the way the word is used in the language. For example, our knowledge that bread is something that is brought about through *baking* is considered a Quale of the word *bread*;

this knowledge is exploited in our understanding of linguistic expressions, such as *fresh bread*, meaning "bread which has been baked recently."

Qualia relations are also referred to as qualia *roles*. The word *role* recalls the notion of *semantic role* used in the domain of verbal semantics to indicate how the various entities associated with a predicate participate in the event expressed by that verb (agents, patients, experiencers and so forth). Qualia roles in GL were first conceived as an argument structure for nouns, and have since been extended to all the major categories.

Qualia encode aspects of a word's meaning that are often attributed as world knowledge by contemporary linguistic theories, i.e., the knowledge we have about objects in the world due to human experience, as in the example of *bread* and *bake* above. In GL, the role of such knowledge is identified when it impacts the behavior of linguistic expressions in usage. We will clarify later how the distinction between lexical meaning and world knowledge is approached in the model.

2.1.1 Qualia and other Formalisms

The Qualia-based system developed in GL to represent lexical meaning is intended to overcome some of the main difficulties encountered by more traditional systems of meaning representation. Traditionally, word meanings have been described in terms of sets of features. The basic idea behind this view is that the meaning of a word is made up of smaller units, called features, components, or primitives (that is, elements that cannot be decomposed any further).

The general strategy to define the meaning of a word in this framework is called lexical decomposition. For example, a *table* may be defined in terms of features such as [inanimate], [concrete], [with legs] and so on. The character of the features that make up a word meaning may vary depending on the category of the word under examination: noun meanings appear to encode features such as [animate], [artifact], [countable], [portable], and [part-of(x)]; verb meanings are assumed to include abstract features such as [act], [cause], [result], [manner], [motion], and so on. Some linguists use a binary notation for these features, such as [+cause] and [-cause]. According to the analysis based on lexical decomposition, a restricted number of features, when appropriately combined, will suffice to define the meaning of all words belonging to a lexicon.

Decomposition into primitives has been very influential in lexical semantics, but it has various shortcomings. The shortcoming that concerns us here is related to the semantic flexibility shown by words; that is, their ability to take on an indefinite variety of possible senses depending on the other words they combine with. This key aspect of word meaning is not easily dealt with in traditional decompositional frameworks. For example, does the verb *like* have two different meanings in "He likes my sister" and "He likes vanilla ice cream"; and if so, how is this difference to be represented in decompositional terms? No traditional decompositional system, no matter how elaborate or refined, is able to answer this simple question in a straightforward fashion.

2.1.2 Decomposition in GL

Although we may say that the GL view of lexical meaning is fundamentally decompositional, in the sense that it is based on the claim that words encode complex concepts that may be decomposed into simpler notions, the general method adopted in GL to define the meaning of words is inverted with respect to the traditional decompositional strategy discussed in 2.1.1. That is, instead of concentrating on how a word meaning may be decomposed, GL examines how a word meaning may or may not compose with other meanings, and how it changes in the different contexts. In other words, GL draws insights about the meaning of a word by looking at the range of its contextual interpretations, and by examining how this range can be predictably derived from the underlying meanings.

For example, with the noun *car*, different aspects of the object are highlighted in the contexts in (2): the car is seen a physical object in (a); as a vehicle in (b); as the part that actually runs and warms up (the car's engine) in (c) and (d) ; as something that that can be locked (the door) in (e); as the parts that produce sound (the car's wheels) in (f). Drawing on such linguistic evidence, GL assumes that all these aspects of *car* (physical object, vehicle, engine, door, wheel) must be considered part of the meaning of the entry *car* in the lexicon, and hence part of its lexical semantics.

(1) a. This car weighs over 2,000 lbs.
    b. We buy vehicles such as cars and buses.
    c. John started the car.
    d. You should warm your car up in winter.
    e. Did you lock the car?
    f. The car screeched down the road.

The procedure adopted in GL to identify the meaning of words requires a system of lexical representation that allows words to change their meaning in different contexts, while maintaining the distinction between word meaning and world knowledge: this is what qualia structure aims to accomplish.

2.1.3 Hidden Events

The first motivation for Qualia relations comes from the analysis of polysemous nominals and adjectives, as described in Pustejovsky and Anick (1988). Specifically, it was proposed that there is a *hidden event* (h-event) in the lexical representation associated with nouns denoting objects that are made for a particular purpose, such as *door* and *book*. A hidden event is defined as the characteristic activity that, when performed, realizes the purpose of the object. Some examples of hidden events for artifactual nouns are given below:

(2) a. a *door* is for "walking through"
b. a *window* is for "seeing through"
c. a *book* is for "reading"
d. a *beer* is for "drinking"
e. a *cake* is for "eating"
f. a *car* is for "driving"
g. a *table* is for "putting things on"
h. a *desk* is for "working on"
i. a *pen* is for "writing with"

According to Pustejovsky and Anick, the reason for including a hidden event in the lexical representation of these nouns is that in certain syntactic contexts this event appears to be present in the interpretation, even though it is not expressed in the syntax. For example, in (4) what is "finished" is the activity of drinking, but this information is not overtly expressed:

(3) They finished the beer. (drinking)

On the other hand, the hidden information is not arbitrary. Rather, it depends on the semantics of noun. For example, in (5) what is finished is the activity of eating, not of drinking.

(4) They finished their cake. (eating)

A similar phenomenon can be observed in Adjective-Noun constructions in (6), where the adjective modifies an activity associated with the noun. For example, a *comfortable chair* is a chair which is good to "sit in", *comfortable shoes* are shoes that are good to "wear" or "walk in", and so forth.

(5) a. a comfortable chair (to sit on)
b. comfortable shoes (to wear, to walk in)

Finally, consider compounds like *dinner dress*, *dessert wine* and *dinner table* in (7). Also in this case, the interpretation entails one (or more) hidden events corresponding to the typical activities associated with the objects *dress* and *wine*: a *dinner dress* is a dress which is "worn" on a special evening occasion, a *dessert wine* is a wine which is "drunk" while "eating" a dessert, and so forth.

(6) a. a dinner dress (wearing)
b. a dessert wine (drinking)
c. the dinner table (eating at)

As we will see, the hidden event introduced in Pustejovsky and Anick corresponds the Telic Quale developed in classic GL.

2.1.4 Qualia Structure

The notion that lexical items can store information relating to hidden events and activities associated with the word is a useful device for helping in the interpretation of linguistic expressions, as noted above. In Pustejovsky (1991), a more elaborated set of relations is proposed, in addition to the hidden event, to represent the meaning of nominals. These relations are called *Qualia*, and the system of relation defining a single concept is called *Qualia Structure*. Qualia Structure consists of four basic roles:

- *Formal*: encoding taxonomic information about the lexical item (the `is-a` relation);
- *Constitutive*: encoding information on the parts and constitution of an object (`part-of` or `made-of` relation);
- *Telic*: encoding information on purpose and function (the *used-for* or `functions-as` relation);
- *Agentive*: encoding information about the origin of the object (the `created-by` relation).

Each Qualia role can be seen as answering a specific question about the object it is associated with:

- *Formal*: What kind of thing is it, what is its nature?
- *Constitutive*: What is it made of, what are its constituents?
- *Telic*: What is it for, how does it function?
- *Agentive*: How did it come into being, what brought it about?

Taken together, the answers to these questions can help elucidate the meanings of words in the language. Adopting the typed feature structure representation introduced in Chapter 1, we can view the qualia structure of a lexical item, α, as the four features below in (8), where F=Formal, C=Constitutive, T=Telic and A=Agentive:

$$(7)\quad \begin{bmatrix} \alpha \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{what } \alpha \textbf{ is} \\ \text{C} = \textbf{what } \alpha \textbf{ is made of} \\ \text{T} = \textbf{function of } \alpha \\ \text{A} = \textbf{origin of } \alpha \end{bmatrix} \end{bmatrix}$$

Recalling the examples of *car* above, we can now identify which qualia role is exploited in the different contexts (9) and represent the results in the feature structures for F and C in (10).

(8) a. This <u>car</u> weighs over 2,000 lbs. (car as material (C))
b. We buy vehicles such as <u>cars</u> and buses. (car as vehicle (F))
c. John started the <u>car</u>. (part of car, engine (C))
d. You should warm your <u>car</u> up in winter. (part of car, engine (C))
e. Did you lock the <u>car</u>? (part of car, door (C))
f. The <u>car</u> screeched down the road. (part of car, wheel (C))

$$(9)\quad \begin{bmatrix} \textit{car} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{vehicle} \\ \text{C} = \{\textbf{engine, door, wheels, ...}\} \end{bmatrix} \end{bmatrix}$$

It should be pointed out that not all lexical items carry a value for each qualia role. Some are left unspecified, while others are populated with more than one value. For example, nouns denoting natural kinds (e.g., *rock*, *fish*, *air*, *sea*) typically do not have a value for the Agentive Quale, since the objects they reference are not products of human creation. While such entities are obviously the product of natural forces, we assume an anthropocentric folk ontology. On the other hand, an artifact such as a *letter* is an entity brought about by a specific activity of "writing", identified through the Agentive role:

$$(10)\quad \begin{bmatrix} \textit{letter} & \\ \text{QUALIA} = & \begin{bmatrix} \text{A} = \textbf{write} \end{bmatrix} \end{bmatrix}$$

The purpose of this same object can be identified through the Telic role, namely, identifying the purpose of the letter as "reading".

$$(11)\quad \begin{bmatrix} \textit{letter} & \\ \text{QUALIA} = & \begin{bmatrix} \text{T} = \textbf{read} \\ \text{A} = \textbf{write} \end{bmatrix} \end{bmatrix}$$

Similar remarks hold for other informational objects such as *book*, *novel*, and *diary*. In later discussions, we will see how to distinguish communicative artifacts such as *letter* from informational artifacts such as book.

Now consider a more elaborate example involving the lexical semantics associated with the noun *house*. First, observe the contexts where different Qualia of the noun are activated (13). Then, look at the proposed GL representation that follows from the observed data (14).

(12) a. He owns a two-story $\underline{\text{house}}$. (house as artifact (F))
b. Lock your $\underline{\text{house}}$ when you leave. (part of house, door (C))
c. We bought a comfortable $\underline{\text{house}}$. (purpose of house (T))
d. The $\underline{\text{house}}$ is finally finished. (origin of house (A))

$$(13)\quad \begin{bmatrix} \textit{house} & \\ \text{QUALIA} = & \begin{bmatrix} \text{F} = \textbf{building} \\ \text{C} = \{\textbf{door, rooms, ...}\} \\ \text{T} = \textbf{live\_in} \\ \text{A} = \textbf{build} \end{bmatrix} \end{bmatrix}$$

2.1.5 Criteria for identifying Qualia Values

In the discussion that follows, we will make a distinction between those qualia values that are lexically specified and those which are introduced by composition in the syntax. Lexically specified (viz., default) values are identified by examining the distribution of the noun in context. For example, in the context of *shatter* in (15a) the noun *window* appears to be used to reference the part represented by the pane of glass, while in the context of *wooden* and *rotting* (15b) it refers to the frame. Because of this evidence, we might consider *pane* and *frame* as default values of the Constitutive quale for the noun *window*.

(14) a. The rock shattered the <u>window</u>.
b. Wooden <u>windows</u> are prone to rotting.

(15) $$\begin{bmatrix} \textit{window} \\ \text{QUALIA} = \begin{bmatrix} \text{C} = \{\textbf{pane, frame, ...}\} \end{bmatrix} \end{bmatrix}$$

Consider now the following contrast:

(16) a. The <u>museum</u> is open until 6:00 pm.
b. *The <u>painting</u> is open until 6:00 pm.

Notice that the noun *museum* in (17a) is being interpreted as the prototypical service event associated with the noun, namely an exhibition. This could be encoded in the Telic quale for *museum*, as shown below.

(17) $$\begin{bmatrix} \textit{museum} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{institution} \\ \text{T} = \textbf{exhibit} \end{bmatrix} \end{bmatrix}$$

However, while the noun *painting* might also be conventionally associated with exhibitions and museums, it does not permit of a similar interpretation in (17b). Such discriminative examples can be systematically used as linguistic evidence for determining what information is lexically associated with the qualia structure of a word. For example, “being exhibited in a museum” would not participate in the direct meaning of *painting* and hence does not constitute a qualia value of the word.

When language accesses the component parts of a word’s meaning with systematic regularity, there is reason to think that those parts might arguably be encoded in the lexical semantics for that word. For example, consider how the noun *car* frequently cooccurs with verbs denoting human actions in subject position: that is, a car can *travel*, *go uphill*, *honk*, *wait for somebody*, and so forth.

(18) a. The <u>car</u> is waiting in the driveway.
b. A <u>car</u> honked from behind.

Such examples are cases of *metonymy*, and are an interesting sense extension from the vehicle to an individual associated with it. Such evidence suggests that the information that a car has a driver is not only part of our world knowledge but is in fact encoded in the lexical entry and available for syntactic selection. In GL, this information is coded as an argument to the predicate that fills the Telic role of car, e.g., *drive*:

(19)
$$\begin{bmatrix} \textit{car} & \\ \text{QUALIA} = & \begin{bmatrix} \text{F} = \textbf{vehicle} \\ \text{T} = \textbf{drive(human,vehicle)} \end{bmatrix} \end{bmatrix}$$

A related metonymic extension is seen with nouns such as *house* and *café*, which are often used to refer to the people who live in or work in the structure.

(20) a. Do you want the whole house waken up?
b. The rest of the house was sleeping.
c. You had the whole café laughing.

Again, such data provide evidence for specific Telic values for these noun concepts; **live_in(human,building)** and **eat_in(human,building)**.

(21)
$$\begin{bmatrix} \textit{house} & \\ \text{QUALIA} = & \begin{bmatrix} \text{F} = \textbf{building} \\ \text{T} = \textbf{live\_in(human, building)} \end{bmatrix} \end{bmatrix}$$

2.1.6 Historical Note

The notion of Qualia in GL originates from the Aristotelian theory of explanation (*aitia*), usually known as the doctrine of four causes. *Aitia* is a Greek term (pl. *aitiai*) meaning 'explanation'. One of the common ways to interpret the Aristotelian scheme is to see it in terms of causal links. According to this interpretation, an *aitia* is the cause of something (x is an *aitia* of y): for example, a sculptor is the *aitia* of a statue, a carpenter is an *aitia* of a table, and so on. In GL, however, Qualia Structure is derived from Moravcsik's (1975) interpretation of the Aristotelian theory of *aitia*. Moravcsik proposes to interpret the theory of *aitia* as a theory of understanding instead of a theory of causation. According to this view, *aitiai* are not simply causes but rather modes of description of an object that allow us to understand it not only by knowing what it is (this is the information provided by the Formal Quale) but also by grasping how it functions (Telic), what its constituents are (Constitutive), and what brings it about (Agentive). GL assumes that the four *aitiai* are relevant not only from an ontological point of view, as for Moravcsik, but also for the characterization of lexical meaning and the modes of compositionality in the grammar.

2.1.7 Qualia in Different Linguistic Phenomena

Explicit reference to qualia structure has proven to be quite useful for representing many linguistic phenomena related to polysemy and lexical ambiguity, which are difficult to deal with using traditional lexical representations. Some of these phenomena, discussed above, include:

(22) a. Contextual modulations of noun meaning, due to the selecting predicate (2); *start/lock* the car

b. Inference of implicit predicates from particular constructions:
Verb-Noun (4, 5) (*finish* the beer/cake);
Adjective-Noun (6) (*comfortable* chair/shoes);
Noun-Noun (7) (*dinner* dress/table).

Additional phenomena include (a) the flexibility of light verbs support verb constructions), and (b) Noun-to-Verb transformations. The first can be seen with verbs such as *make*, *take*, and *have*, where specific information in the Qualia Structure of the complement is exploited in the overall interpretation of the construction, as illustrated in (24), where *take* is interpreted as *ingest* in the context of *tablet* and as *use to travel* in the context of *train*;

(23) Light verbs specifications:
i. Take a *tablet* (TELIC = ingest)
ii. Take a *train* (TELIC = travel with)

Noun-to-Verb transformations, on the other hand, involve a category shift of a noun to a verb, in which the noun's Telic role becomes the verb meaning. Examples include the verbal forms of the nouns *fax*, *microwave*, and *lace*, as demonstrated below in (25).

(24) Noun-to-Verb transformations:
a. *fax* a document: (TELIC = send)
b. *microwave* the chicken: (TELIC = cook)
c. *lace* the shoes: (TELIC = tie)

There are many phenomena other than polysemy that receive a natural interpretation if Qualia are assumed as one of the basic components of lexical meaning. For example, the apparent necessity of adjuncts in constructions such as short passives (26), middles (27) and past participle constructions (28) may be analyzed as determined by the fact that the noun and the predicate in the construction form a *qualia pair*, that is, a combination in which the predicate expresses one of the qualia values of the noun (like *picture-paint*, *book- read*, or *house-build*), a phenomenon called *co-specification* in GL (cf. Chapter 5). In GL terms, qualia pairs are considered uninformative in a typical discourse context, since the verb predicates information which is already encoded as part of the noun's meaning. Typically, this is associated with properties of definiteness of the NP whose head is part of the qualia pair. Hence, this is why an adjunct is needed to avoid the sense of uninformativeness in these uses, as illustrated in (26) below. Similar remarks hold for short middles (27) and adjectival uses of past participles in Adj-N pairs (28).

(25) Short passives (TELIC(picture) = paint):
  a. *This picture was painted.
  b. This picture was painted *in 1604*.

(26) Middles (TELIC(book) = read):
  a. *This book reads.
  b. This book reads *easily*.

(27) Adjectival Use of Past Participles (TELIC(book) = read):
  a. *a built house;
  b. a *recently* built house.

Throughout the book, in fact, we will see how qualia structure is implicated in far more grammatical constructions and licensing operations than one would at first imagine.

2.2 The Four Qualia Roles

In the following sections, we discuss in more detail the properties of the individual Qualia introduced in the previous section. For now, we restrict our discussion to the semantics of nouns. For each role, we give a definition and provide examples that clarify how it is interpreted and how it can be used in lexical analysis. The illustration of Qualia Theory we present is based on theoretical generalizations formed from the empirical analysis of lexical distributions in language.

2.2.1 Formal

A lexical item carries information about the basic conceptual category with which it is associated. In GL, this information is coded in the Formal quale (FORMAL) of the word. Particularly, the Formal quale establishes a relation between the entity denoted by a word (e.g., *dog*) and the category it belongs to (i.e., *animal*). It is this relation that enables one to grasp the nature of an entity by distinguishing it among other kinds. It answers the question: "What kind of entity is x?" For example, a *rock* is a natural kind, a *table* is an artifact, a *car* is a vehicle, a *park* is a location, *water* is a liquid, a *plant* is a living thing, a *fish* is an animal, a *hand* is body part, a *glass* is a container and so on. Sometimes, more classifications are possible for the same type of object: for example, a *knife* can be a weapon or a kitchenware. Moreover, classifications at different levels of generalization are available for reference: for example, in (29) *water* is seen as a liquid (its immediate superordinate), a fluid (a higher superordinate) or a substance (the highest superordinate). These are seen in the *such as*-construction, shown below.[1]

(28) a. a liquid such as water;
b. fluids such as water or air;
c. substances such as fluids, salts, glucose and carbon dioxide.

The hierarchy reflected by these constructions is illustrated in 2.1 below.

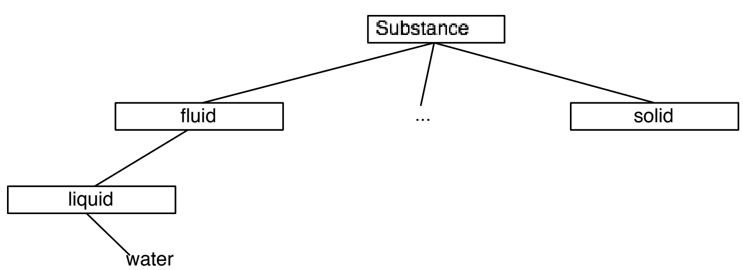

Figure 2.1: Lexical Hierarchy reflecting Type Structure

[1] This construction has been used for lexical acquisition purposes in Computational Linguistics research (cf. Hearst (1992,1995), Lin (1998)), along with other grammatical patterns indicating specific function or semantic roles (cf. Pustejovsky et al, 1993, Wilks et al, 1996).

As another illustration, consider the representation of the Qualia structure of *rock* in (30), focusing on the Formal role, and the accompanying linguistic distributional evidence for such a classification in (31) . Note the the examples in (31) give us information not only about the superordinate category (natural) but also about properties of rock: for example (31c) tells us that rocks are made up of hard material, a point we will come back below.

(29) $\begin{bmatrix} \textit{rock} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{physical} \end{bmatrix} \end{bmatrix}$

(30) a. inanimate objects such as rocks;
b. natural elements such as rocks, soil, timber and so on;
c. hard material such as rocks.

Categories associated with words may be analyzed as clusters of characteristic properties. Salient properties of a category (i.e., the properties which enter into the constitution of the concept for that category) answer the question, "What makes this y an x?": for example, "What makes this animal *a dog*?", "What makes this thing *a table*?" and so forth. We may identify salient properties for an entity by looking at the distribution of the noun that references it. For example, the example of *rock* in (31c) illustrates that rocks are conventionally classified as made of hard material, despite the fact that a distinction between hard rock and soft rock may be identified at a more technical level (marble is a hard rock, while chalk is a soft rock). This information, however, pertains to the Constitutive role, as it involves the material which rocks are made of. This will be discussed in the next section.

We will distinguish two kinds of properties that an object may have, for purposes of classifying this object relative to other entities:

(31) a. Properties that are relevant for the physical taxonomic classification of an entity;
b. Properties describing an entity's internal and external constitution, typical behavior, use, purposes, or function.

The Formal role typically provides access only to the properties of an object that are relevant for its taxonomic classification. These are the properties that a word inherits from its superordinate. It is on the basis of these properties that we are able to draw class inferences. For example, by virtue of its Formal relation with liquid, water inherits the defining properties of liquids, such as fluidity. Note that the constituents of water, namely hydrogen and oxygen, do not inherit the properties of water, since the relation they hold with water is not a Formal relation (is-a relation) but a Constitutive relation (made-of relation). For now, we will assume that property inheritance (class inference) only occurs along with Formal qualia values.

Some lexical items appear to have more than one Formal role value. These are cases of *multiple inherence* over the Formal quale: for example, the noun *book* refers to both a physical object (as in *hand me the book*) and an informational object (as in *believe the book*); the noun *lunch* refers both to an event (as in *during lunch*) and the food eaten (as in *a spicy lunch*); the noun *house* denotes both a physical object (as in *build a house*) and a location (as in *leave the house*). In the GL model, these are called *dot objects* and will be introduced in Chapter 3.

We can summarize the properties of the Formal Quale as providing access to three kinds of information:

(32) a. The basic category associated with the word (i.e., its semantic type);
b. The position of the word in the hierarchy of types following from this association;
c. The salient properties which enter into the definition of the type, which are inherited by the word along the Formal role.

Pustejovsky (1991,1995) provide a list of Formal attributes (or *factors*) for the class of nouns denoting concrete entities. These include physical characteristics such as the following:

(33) a. Spatial characteristics, intrinsic orientation;
b. Size and dimensional properties;
c. Shape and form;
d. Color.
e. Position.

Each attribute may be filled by a value: for example, in *a red car*, *red* is the value (filler, descriptor) of the Formal attribute "color"; in *a long dress*, *long* is the value of the Formal attribute "dimension"; in *a round table*, *round* is the value of the Formal attribute "shape", and so forth. Some combinations may be ambiguous: for example, a *red pen* might mean "a pen that is red" or "a pen that writes in red ink". In the first case, we assume that the adjective activates the Formal attribute "color" of the noun, while in the second case, it appears that it is the Constitutive or the Telic which play a role in the interpretation. Obviously, several modifications, referring to different factors, might be present in the structure of a single NP, such as in (35) below, where FACTOR_F refers to a named factor of the Formal role:

(34) a *large* (Size_F) *round* (Shape_F) table

Conversely, a named factor value can generally be filled only once for each modified noun.

(35) *a *round* (Shape_F) and *square* (Shape_F) table

Lexical meaning often provides default values for the different Formal factors or attributes. Default values are inherent properties of entities which enter into their denotation. For example, we might assume that the "Size" value associated with the noun *ant* is *small*, when evaluated relative to the superordinate class for the noun *insect*. This can be supported from such uses as seen in (37):

(36) ... an *insect* as small as an *ant* ...

Default values, however, may be updated from discourse context in composition: for example, in *large ant*, context makes us update the value of the Size factor from *small* (default) to *large (for an ant)*. Comparison classes of adjectives are suggested by specific information from the Formal role; e.g., in *a large ant* vs. *a small dog*, we have basic category and ontological classification information specified that gives us a way to constrain the interpretation of relative interpretations of "Size" in such constructions.

It is often the case that there is no value specified lexically for a particular Formal factor. For example, one might assume that part of our knowledge of ants is that they are typically black in color. Modeling such information lexically entails specifying a default value for the color factor associated with the Formal role associated with the lexical item *ant*. For many natural kinds, however, no such default color can be associated with the individuals associated with that kind, e.g., dogs, cats, rocks. In such cases, the value for a particular Formal factor may be introduced in context. For example, in *black dog*, the expression compositionally introduces the value *black* for the color factor associated with the Formal for the noun *dog*.

As mentioned above, the Formal role can be seen as a containment relation: if the Formal(water) = liquid, then the relational form can be read as `is-a`(water, liquid). The two terms participating in the Formal relation tend not to cooccur in the same local context since they are in a paradigmatic relation with each other. We may, however, occasionally find both terms in constructions such as those shown in (38).

(37) FORMAL–specific Constructions:
a. NP *such as* NP: *events such as lectures, walks, tours and meetings*;
b. *such* NP *as* NP: *such areas as children's playground*;
c. NP *and other* NP: *rum and other spirits;*
d. NP *or other* NP: *insects or other animals*
e. NP, *including* NP: *recyclable materials including glass*;
f. NP, *especially* NP: *cool temperate countries especially Europe and North America*;
g. *favorite* NP is NP: *Mario's favorite food is pasta*.

On the other hand, one regularly finds Formal factors expressed through adjectival, nominal, verbal or prepositional descriptors, as illustrated in (39-42):

(38) a. a *flat* screen (Shape$_F$)
b. a *leaning* tower (Shape$_F$)
c. a *straight* line (Shape$_F$)
d. a *thick* sweater (Dimension$_F$)
e. an *expensive* car (Cost$_F$).

(39) a. the *height* of the wall (Dimension$_F$)
b. the *length* of the table (Dimension$_F$)
c. the *depth* of the sea (Dimension$_F$).

(40) a. the *front* of the house (Orientation$_F$)
b. the top of the table (Orientation$_F$)
c. the *foot* of the stairs (Orientation$_F$)
d. the *head* of the bed (Orientation$_F$)
e. the *facade* of the building (Orientation$_F$);

(41) a. *clean* the table (Surface$_F$)
b. *wipe* the floor (Surface$_F$)
c. *paint* the wall (Surface$_F$).

In this section, we have focused mostly on the basic properties of the Formal role for lexical items denoting natural kinds. In later chapters, we will examine how the Formal is defined for words denoting more complex entities, including objects with specific functions and purposes (artifactuals), objects which inherit from multiple superordinates (dot objects), as well as relational nouns.

2.2.2 Constitutive

In the previous section, we saw how the representation of word meaning encodes the basic categorical knowledge associated with a lexical item through the Formal role of the qualia structure. The classification resulting from this information is essential for differentiating objects within a class. To make this clear, consider how we analyzed the noun *water*. We said *water* was defined in terms of its relation to the supertype, *liquid*, which in turn was defined as a *fluid*, and so on, through the Formal role. Similarly, the noun *woman* has a Formal value of *human*, which in turn has a Formal value of *animal*, and so on.

Upon closer examination, however, it becomes clear that the values associated with the Formal roles for these two cases are quite different: namely, there is a fundamental distinction between inherently individuated *things*, such as humans, tigers, and trees, and inherently undifferentiated *stuff*, such as water, air, and sand. This is in fact known as the "count/mass" distinction in linguistics, and is a key aspect in determing the grammatical behavior and distribution of nouns in language. Because this distinction is so central to the conceptualization of our world, it plays a central role in how the Formal role for nouns is interpreted in a lexicon.

In GL terms, there is not actually a mass/count distinction in the grammar. Rather, this distinction is an artifact of how the Formal role interacts with the *Constitutive* role for a lexical item. To understand how, we look at a range of examples somewhat informally. We can think of the Constitutive (CONST) role associated with an object as encoding information about what is "inside" that object, the material the object is made of (i.e., its stuff), and the parts it consists of.

For example, the noun *car* denotes a vehicle (its Formal role value), but is made of many different part, e.g., chassis, engine, seats, and so on. We say that these parts denote the value of the Constitutive of the noun *car*. Similarly, a tree has many parts, such as a trunk, branches, leaves, and roots. So, while the noun *tree* denotes a plant (its Formal role value), these aspects would designate the value of the Constitutive role for *tree*.

Let us now return to mass terms such as *water* and *air*. One of the defining characteristics of a mass term is that smaller parts of the material being referred to are still that material. In other words, any decomposition results in the same value. This has a very elegant representation in GL, where count and mass terms are distinguished by their qualia structures. We will make the following distinction here. There is formally one relation of $part\_of(x, y)$, but when the types of x and y are identical, $x : \tau = y : \tau$, then we essentially have the $made\_of$ relation. When the types between x and y are distinct, however, $x : \tau = y : \sigma, \tau \neq \sigma$, then we have a $part\_of$ relation between heterogeneous objects. This is shown in (43) and (44) below.[2]

(42) COUNT NOUN: (where $\alpha \neq \beta$)

$$\begin{bmatrix} \textbf{N} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \alpha \\ \text{C} = \beta \end{bmatrix} \end{bmatrix}$$

(43) MASS NOUN:

$$\begin{bmatrix} \textbf{N} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \alpha \\ \text{C} = \alpha \end{bmatrix} \end{bmatrix}$$

Count nouns have parts and possibly made of relation, and mass terms have no distinction between made of and parts.
The constraint in (44) is known as the *Formal-Constitutive Equivalence Constraint*. Another representation for the qualia structure for mass terms that we will use is: $[\text{F/C} = \alpha]$.

Now let us revisit the qualia structure for the noun *water* from the previous section. As a mass noun, there is no distinction made between the Formal and Const qualia, and we will represent it as shown in (45).

(44)
$$\begin{bmatrix} \textit{water} \\ \text{QUALIA} = \begin{bmatrix} \text{F/C} = \textbf{liquid} \end{bmatrix} \end{bmatrix}$$

Notice that this ensures that any type-based inference through the Formal role is still available to words denoting concepts pertaining to "stuff": i.e., water is a liquid, liquid is a fluid, and so on.

The role played by the Constitutive quale would be significant enough from a semantic point of view if this were its only contribution to linguistic modeling. This is, however, only a small part of the descriptive power of this role. Since most nouns in the lexicon are count nouns (at least in English), the Formal-Constitutive Equivalence Constraint will not be relevant.

[2] In later chapters, we will use a formal device called *reentrancy*, which allows multiple feature labels (e.g., qualia roles) to point to the same value, e.g., α.

This means that the Constitutive is available for encoding concepts associated with a broad range of *part-of* relations. That is, if "x has a Constitutive value of y, then $part\_of(y,x)$."

For example, for the noun *car*, the Constitutive expresses the relation of the car to its constituent elements, e.g., chassis, engine, windows, seats, and so on.

(45) $\begin{bmatrix} \textit{car} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{vehicle} \\ \text{C} = \{\textbf{chassis, engine, seats, windows, ...}\} \end{bmatrix} \end{bmatrix}$

Besides expressing the internal makeup of an object, the Constitutive may also introduce the relation between an object and the larger object it is is logically a part of, if such a dependency exists[3]. The noun *roof*, for example, denotes part of a larger entity, namely the concept of a *building* This is characterized by reversing the direction of the Constitutive to allow reference to entities which are themselves parts of things. We will call this relation the *Inverse Constitutive (C_I)*. Hence, if "x has an Inverse Constitutive value of y, then $part\_of(x,y)$."

(46) $\begin{bmatrix} \textit{roof} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{phys} \\ \text{C}_I = \textbf{building} \end{bmatrix} \end{bmatrix}$

Notice that both of these aspects of the Constitutive role come together in the semantics of a noun such as *room*, which is both part of something (a building) and has parts in it (walls, floor, ceiling). This can be expressed in the lexical representation shown below in (48).

(47) $\begin{bmatrix} \textit{room} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{space} \\ \begin{bmatrix} \text{C} = \{\textbf{walls, floor, ceiling, ...}\} \\ \text{C}_I = \textbf{building} \end{bmatrix} \end{bmatrix} \end{bmatrix}$

The linguistic relevance of the Constitutive quale for the noun *room* can be seen in (49), where in the context of *paint*, *room* is able to assume the interpretation of "the walls from the room", while in the context of *sweep* it refers to the "floor of the room."

(48) a. John was going to paint his room ([CONST = walls]).
b. She has swept the room ([CONST = floor]).

[3]This is explored in Pustejovsky and Anick (1988) and Vikner and Hansen (1994).

Similarly, in the selectional context of the verbs *cut* and *sharp*, the noun *knife* can be used to denote specifically the part that enables the cutting, i.e., the blade (50).

(49) a. The knife cut his finger ([CONST = blade]).
b. a sharp knife ([CONST = blade])

(50) $$\begin{bmatrix} \textit{knife} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{phys} \\ \text{C} = \{\textbf{blade, handle}\} \end{bmatrix} \end{bmatrix}$$

Defining what counts as a part of an entity is a controversial topic. The matter is hotly debated especially within the theory of mereology (the study of parthood relations). In a way, the word *part* may be used to indicate any portion of a given entity, regardless of whether, for example, that portion is attached to the rest of the object, as with "the *handle* of a door", or undetached, as "the *cap* of a pen." However, among logicians, philosophers and ontologists it is often assumed that the "legitimate" parts of an object are those that demonstrate the following characteristics:

(51) a. Parts are available in discourse as individual units;
b. Parts make a functional contribution to the entity;
c. Parts are cognitively salient.

The last point is meant to rule out arbitrarily demarcated regions or portions of an object, such as *the lower part of the wall*, which is clearly part of the wall, but not to be identified as a legitimate "part" of the wall, since it refers to a part of the wall with no clear boundaries. It is also generally assumed that a well-formed part-of relation should consist of elements of the same general type; a part of some physical entity will also be a physical entity, part of some time period will also be a time period, and so forth. According to this view then, the "weight" of a body does not figure among its parts, even though it will constitute an essential attribute (or factor) of that body.

The GL approach to the parthood relation expressed by the Constitutive and Inverse Constitutive qualia is driven by the grammatical reflexes that the constituency of an entity expresses, in the way we reference that entity in language. That is, in GL it is assumed that the Constitutive quale specifies only those parts of an entity that are relevant for the linguistic behavior of the noun expressing the whole, as the *walls* of a room. On this view, entities may have several parts, only some of which may be recognized as values of the Constitutive quale for the word denoting that entity. Other parts may not be lexically encoded as part of the qualia structure of the word, even though they are semantically implied as parts of the entity. Therefore, it will be useful to distinguish between the following two noun types:

(52) a. Nouns that specify a value for CONST lexically;
b. Nouns that lexically lack a value for CONST.

Typically, nouns of type (53a) are words that lexical whole objects that allow for reference to parts, such as *room* or *knife*. These nouns may express one of the default values of CONST syntactically, as for example in a genitive construction (54a), or in a construction where the first term introduces reference to a part of the entity denoted by the second term (54b):

(53) a John was going to paint the room's walls.
b. John was going to paint the walls of the room.

The behavior observed above follows from the fact that the noun *wall* is lexically specified within the qualia structure of the noun *room*. Notice however, that there are many ways to dynamically elaborate parts of objects that are not lexically specified in the noun. For example, although most rooms would arguably have corners as spatially defined partitions within the region defined by the room, there is no justification for calling a corner a "part of" the room, assuming the methodology we have adopted here. Consider the sentence in (55), where a part of the room has been identified compositionally in context:

(54) He was standing in the corner of the room.

Although the most conventional choice for the Constitutive relation is the part-of relation, it can also be used in a broader sense to designate the relation of material constitution (made-of relation). This interpretation dates back to the Aristotelian scheme, where, for example, marble is described as being in a constitutive relation to the statue it is part of. In this sense, nouns such as *river*, *lake*, and *sea* have CONST values of water, since they are largely composed of this material. The qualia structure for the noun *river* encodes this value in its CONST role directly, as illustrated below, where we assume the FORMAL value for *river* is region or space, filled with the material designated by the CONST value.

(55) $\begin{bmatrix} \textit{river} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{space} \\ \text{C} = \textbf{water} \end{bmatrix} \end{bmatrix}$

While the FORMAL role can be modified by spatial predicates (cf. (57)),

(56) a. They crossed the river. ([FORMAL = space])
b. The river is wide. ([FORMAL = space])

in fact, the constitutive value is directly referenced by the verbs *freeze* and *pollute* as shown below.

(57) a. The river had frozen during the severe weather. ([CONST = water])
b. The river became polluted. ([CONST = water])
c. the banks of a polluted river ([CONST = water])

While the example in (58b) exploits the qualia information associated with the CONST value of the noun *river* in order to arrive at the interpretation of "water being polluted", CONST values of a noun can be explicitly expressed in syntax, as demonstrated in (59).

(58) a. The water in the river is polluted.
b. polluted river water.

Another typical way of expressing the made-of relation in English as well as in other languages is through nominal compounding. In this construction, the value for CONST is generally introduced in composition as a modifying bare noun. For example, in the Noun-Noun compounds in (60), the first noun (*plastic*, *paper*, *leather*) expresses the value of the CONST quale for the head of the expression: e.g., a plastic bag is a bag made of plastic, a paper cup is a cup made of paper and so on.

(59) a. plastic bag
b. paper cup
c. leather shoes
d. gold watch
e. milk chocolate

(60) $\begin{bmatrix} \textit{plastic bag} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{bag} \\ \text{C} = \textbf{plastic} \end{bmatrix} \end{bmatrix}$

The made-of relation may also be introduced in adjectival modification constructions, such as in the Adjective-Noun constructions in (62):

(61) a. a golden ring;
b. a wooden floor;
c. a metallic paint.

Finally, the made-of relation may be introduced indirectly, i.e., by referencing specific attributes related to the material. For example, with the expression *heavy chain*, we refer to the property "weight" that inheres of the object by virtue of its constitutive material; a chain may be *light* or *heavy* depending on the material it is made of. Therefore, we say that the word *heavy* introduces the value "heavy" to the attribute (or factor) "weight" of the CONST in the Qualia Structure of *heavy chain*. This is represented below, where the adjective *heavy* introduces the material of which it is measuring the weight.

(62) $$\begin{bmatrix} \textit{heavy chain} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{chain} \\ \text{C} = \textbf{heavy(material)} \end{bmatrix} \end{bmatrix}$$

Similar remarks hold for other substance modifying attributes, such as "temperature". Consider, for example, the expression *cold water*, where *cold* modifies the factor "temperature". Because *water* is a mass term, the FORMAL value is identical to its CONST value. For a mass noun such as *soup* in he expression *thick soup*, the adjective *thick* is modifying a property of the density or viscosity of the liquid.

To conclude this section, we review some of the constructions seen in language that are specific indicators of constitutive relationships. These are shown in (64), where $Const(x, y)$ indicates the relation "x is the CONST of y".

(63) CONST–specific Constructions:
a. N_1's N_2: $Const(N_2, N_1)$
the room's wall;
b. N_2 of N_1: $Const(N_2, N_1)$
the door of the car;
c. NP_2 is a part of NP_2: $Const(NP_2, NP_1)$
brain is a very sensitive part of the body;
d. NP_1 made of NP_2: $Const(NP_2, NP_1)$
monuments made of stone and marble

e. NP_1 of NP_2:
house of wood; $Const(N_2, N_1)$
f. NP_1 consists of NP_2: $Const(NP_2, NP_1)$
the orchestra consists of ninety performers;
g. NP_1 containing NP_2: $Const(NP_2, NP_1)$
a forest containing dead trees.

2.2.3 Telic

In this section, we turn to the analysis of the Telic Quale. The Telic Quale, TELIC (from the Greek term $\tau\epsilon\lambda o\varsigma$, meaning "end" or "goal"), encodes information about the intended use or function of an object. Specifically, it expresses the relation that allows us to grasp what an entity is by knowing what it is used for, and what we normally do with it. Therefore, the value expressed by the Telic relation of a noun often corresponds to the activity in which the object named by the noun is typically involved; for example, a food item like *cake* has the Telic value of eat, while an instrument for writing such as *pen* has a Telic value of $write\_with$.

(64) $$\begin{bmatrix} \textit{cake} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{food} \\ \text{T} = \textbf{eat} \end{bmatrix} \end{bmatrix}$$

(65) $$\begin{bmatrix} \textit{pen} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{tool} \\ \text{T} = \textbf{write\_with} \end{bmatrix} \end{bmatrix}$$

The activity specified in the Telic should be interpreted as a potential activity or characteristic property of the object; as such, it is modally subordinated to the specific context which enables the possibility of this function to be activated. For example, a *pen*, has the characteristic of being used for writing, even if this characteristic is not exploited in a specific context. In similar fashion, *mail* is for sending even if it is not sent at all. This characteristic is a persistent property of the object, i.e., a property that persists through time. What is important for the Telic relation, however, is that when the activity is performed, the main purpose of the object is satisfied. We will examine in more detail this aspect of the Telic quale in the section dedicated to the relation between Telic and Modality in Chapter 5.

There is a good deal of empirical evidence suggesting that certain nouns encode information about the intended use or function of the object to which they refer (see also the examples in section 2.1.3). For example the contexts in (67) clearly evoke a writing event made available by the noun *pen* even if this event is not expressed in the syntax. That is, the contexts in (67) are meaningful only if a writing event is assumed in the semantic representation of the sentence.

(66) a. This pen does not work well. (does not write)
b. Can I use your pen? (for writing)
c. Have you got a red pen? (ambiguous, which writes in red)

Similarly, an eating event is made available by the noun *cake* in all the contexts in (68).

(67) a. Any chocolate? Not after that cake! (after eating)
b. I prefer cake to biscuits. (prefer eating)
c. We skipped the cake and settled for another coffee. (skipped eating)

At times, the implicit predicate does not satisfy the intended function of the object but expresses one of the activities associated with its purpose. For example in the contexts in (69), the implicit predicate appears to be the "departure" of the train, while the relevant Telic value associated with a train would be that of *travel*.

(68) a. There's no train till 7:00 pm. (there is no departing)
b. The train was delayed for an hour. (the departure was delayed)
c. I left in time to catch the early train. (departing early)

In GL, the examples in (69) do exploit information relating to the function of the noun *train*, but it is more indirect than in our previous examples. In these cases, the elided predicate seems to relate to the *departure* of the train rather than simply *traveling* on it. Such knowledge is not directly part of the qualia structure per se, but is derived from it, by means of *projective* operations over these values. In other words, the act of "departing" is a projective activity, and is available for interpretation as one of the components of the overall Telic activity of "transportation". These projected activities are treated as conventionalized attributes in GL. The representation of conventionalized attributes is discussed in detail in section 2.3, while a discussion of the mechanism of projection is presented in Chapter 5.

Another important piece of evidence for the Telic relation comes from Adjective-Noun constructions where the meaning of the adjective provides information regarding the intended use and function of the noun's referent.

For example, in the context of *customer* in (70) *next* is referencing an ordering over the servicing of individual patrons (the customers), i.e, 'next to be taken care of', while in the context of *slide* it references the order of the viewing or presentation of the individual slide, i.e., 'next to be projected'.

(69) a. the next customer (to be taken care of)
c. the next slide (to be projected)

Within GL, it is assumed that the contextual variation in meaning of *next* in (70) is determined by the fact that the adjective targets a sub-portion of the semantics of the head noun, particularly its Telic relation. The specific mechanism allowing this modification will be detailed in Chapter 5. Notice that context can be an important determining factor with such constructions. Imagine a conversation in a moving vehicle, where the passenger tells the driver to "turn left at the next intersection." In this case, *next* refers to the ordering imposed by the path assumed by the moving car.

Further examples in which the Telic relation of a noun plays a role in the interpretation of the modifying adjective (in both attributive and predicative position) are given below:

(70) a. This is a difficult problem (to solve).
b. This is a difficult question (to answer).

The Telic of *problem* can be said to involve its solution, while the Telic of *question* involves its answer. The adjective *difficult* modifies the Telic aspect of the noun. Similar behavior is seen with adjectives such as *ready* in (72).

(71) a. Your coffee is ready (to drink).
b. There's some lunch ready in the kitchen (to eat).
c. The car is ready (to drive).

In (73) we give a list of adjectives selecting the information made available by the Telic quale of the noun. For each adjective, we give a contextualized example, and indicate in brackets the telic value which is being modified.

(72) Telic selectors:
fast food (to eat), a slow oven (to cook), a short novel (to read), a complex question (to answer), an easy place (to get to), useful, an effective antibiotic (to cure), agreeable, avoidable costs (to pay), enjoyable, a good doctor (to heal), a bad singer (to listen to), an interesting book (to read), ready meals (to eat).

Implicit predicates in Adjective-Noun constructions may not match the Telic information of the noun while at the same time are clearly associated with it. For example in (74) the adjective *heavy* appears to modify the activity of *digesting* instead of the activity of *eating*. This is shown by the fact that we can paraphrase *heavy* with 'hard to digest' rather than 'hard to eat'.

(73) Heavy foods such as dairy products and meat.

In GL terms, this is again a case of exploitation of one of the *conventionalized attributes* associated with the Telic, which we will examine in more detail shortly.

Finally, as pointed out in 2.1.3, the Telic relation appears to play a role in the interpretation of nominal compounds. For example, in the compounds in (76) the Telic information encoded in the head nouns (e.g., *shop*, *glass*, *bus*, *train*) makes available the semantic relation between head and non-head that is exploited in the interpretation of the compound: a *book shop* is a shop where book are "sold", a *wine glass* is a wine for "holding" wine, a *school bus* is a bus for "transporting" kids to school, and so forth.

(74) a. a book shop (selling)
b. a wine glass (holding)
c. a school bus (transporting)
d. a freight train (transporting)

In many constructions, the Telic information encoded in the non-head noun interacts with that of the head and constrains the interpretation. Compare *tooth brush* (cleaning) vs. *hair brush* (combing) in (76).

(75) a. a tooth brush (cleaning)
b. a hair brush (combing)

The information expressed by the Telic is characteristic of artifactual objects, i.e., objects created for a particular purpose. However, this information appears to be present also in nouns denoting a variety of concepts:

(76) a. functional locations: *library*, *gym*, *church*, *school*;
b. professions: *doctor*, *teacher*, *lawyer*;
c. agentive nominals (individuals engaged in an activity, either habitually or occasionally): *runner*, *passenger*, *movie goer*.

Functional locations, for example, are locations that are designed for the performance of particular activities. These activities may be encoded as Telic information of the corresponding nouns and may be activated in particular contexts. For example, in the context of *attend* in (78), *church* references the activity specified in its Telic quale and is reinterpreted as "church service" contextually, i.e., the event of the mass.

(77) He no longer attends church. (mass)

The Telic value of a noun denoting an artifact may be populated with more than one item or it may have none. In other words, not all nouns denoting artifacts encode Telic information. According to Asher and Pustejovsky (2006), this is because not all functions that we ascribe to artifacts "make it over to the lexicon". That is, while all artifacts by definition have a characteristic function corresponding to the purpose for which they are built, information about this function is not necessarily incorporated in the meaning of the corresponding lexical item. We will come back on this issue, which touches on the relationship between the conventional aspects of word meaning and the general knowledge we associate with objects, in Chapter 3.3.

There are two main types of Telic, as discussed in Pustejovsky (1995): Direct Telic and Purpose (or Indirect) Telic. With the Direct Telic, the entity denoted by the noun is realized as the object of the activity (e.g., *cake* is the object of the Telic value, eat), while with the Purpose Telic, it is realized as subject or indirect object (e.g., *pen* is the subject or indirect object of the Telic value $write$). The Direct Telic characterizes the entity as something which one acts on directly. The Indirect Telic characterizes the entity as something which is used for facilitating a particular activity or has the function of carrying it out. Within the Indirect Telic, Busa et al. (2001) distinguish two subtypes: (i) the entity is an instrument (e.g., *pen* relative to $write$); (ii) the entity is an agentive noun (e.g., *singer* relative to $sing$). This information may be encoded in the lexical representation of *cake*, *pen* and *singer* as follows:

(78) $$\begin{bmatrix} \textit{cake} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{food} \\ \text{T} = \textbf{eat(human,food)} \end{bmatrix} \end{bmatrix}$$

(79) $$\begin{bmatrix} \textit{pen} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{tool} \\ \text{T} = \textbf{write\_with} \end{bmatrix} \end{bmatrix}$$

(80) $\begin{bmatrix} \textit{singer} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{human} \\ \text{T} = \textbf{sing(human, song)} \end{bmatrix} \end{bmatrix}$

There are a number of Qualia-specific constructions for the Telic, that is, constructions where the two terms of a Telic relation co-occur. When the Telic activity being expressed corresponds to the Telic value specified in the noun (i.e., to the default Telic), the resulting expression is what is called a qualia pair. A qualia pair may take the form of a verb-noun pairing (82a), an Adjective-Noun pairing (82b), or a compound (82c):

(81) a. The child drank some water.
b. drinkable water;
c. drinking water.

Notice that *-able* adjectives appear to impose a specific interpretation on the Telic activity of the noun they modify. For example *drinkable* does not simply mean 'that can be drunk' but rather 'that is good for drinking'. Other examples of constructions where *-able* adjectives expressing the Telic value of the head noun are given below:

(82) a. The rent is payable monthly.
b. a very readable text-book;
c. a very playable game.

When the Telic activity being expressed does not correspond to the Telic value specified in the noun, we say that the expression updates the Telic information associated with the noun in composition. Consider, for example, the complex nominal *shopping bag*. The lexically specified Telic value for *bag* is *hold*, by virtue of it being a container; this value is updated in composition by the modifier *shopping*, which becomes the Telic value of the overall expression.

(83) $\begin{bmatrix} \textit{shopping bag} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{container} \\ \text{T} = \textbf{shopping} \end{bmatrix} \end{bmatrix}$

Similar remarks hold for the compounds listed below.

(84) a. ironing board (used for ironing)
b. swimming pool (used for swimming)

c. dining room (used for dining)
d. frying pan (used for frying)
e. cutting knife (used for cutting)

Other specific constructions expressing Telic values (either lexically specified or updated compositionally) are given below in (86), where $Telic(x, y)$ indicates the relation "x is the TELIC of y".

(85) TELIC-specific Constructions:
a. an NP to V: $Telic(V, NP)$
a book to read;
b. an NP worth V-ing: $Telic(V, NP)$
a question worth asking;
c. the NP merits/deserves V-ing: $Telic(V, NP)$
This book deserves reading.;
d. enjoy/prefer V-ing NP: $Telic(V, NP)$
enjoy listening to music / prefer watching television;
e. an Adj NP to V: $Telic(V, NP)$
a difficult question to ask;
f. an NP (used) for V-ing: $Telic(V, NP)$
a spade (used) for digging;

Natural Telic

While the Telic of an artifact gives information about the intentional activities that satisfy the object's design or purpose (*pens* are for *writing*), the Telic of a natural kind (*human*, *dog*, *water*, and so forth) encodes information about the actions and properties that the object engages in, but that are not in any way intentional or purposive. For example, when we describe the analytic properties associated with humans, as expressed in (87), we are predicating a "natural Telic" property of that entity.

(86) a. Humans breathe/think.
b. Rivers flow.
c. The heart pumps blood.

It is not the intentional purpose of a heart to pump blood, but it is a necessary activity for the object so defined. Likewise, a river does not intentionally flow, but this is a necessary property of a body of water if it is to qualify as a river.

The linguistic relevance of the property of "flowing" for the noun *river* appears in (88), where in the context of *swift*, *lazy* etc., *river* allows for reference to its Natural Telic value: in other words, a *swift river* is a river 'which flows at high speed'. Additional examples modifying the attribute of "speed of flow" are shown below.

(87) a fast/rapid/slow/lazy river (flowing)

Similarly, corpus analysis shows that the expression *flowing river* is uncommon while expressions specifying the manner of flowing are totally normal and frequent, as seen in (89).

(88) a. a fast/quietly/slowly/steadily flowing river;
b. a gently/peacefully flowing river.

This evidence supports the GL view that *flowing* is part of the meaning of *river*, and hence is uninformative as a modifier to the noun *river*. By contrast, in the expressions in (89) the modifier introduces new conceptual material, thus satisfying the requirement that every linguistic expression must be informative in the discourse context (see 2.1.7).

Given such behavior, we now return to the qualia structure for the noun *river*, and supplement the Formal and Constitutive values with a specification for the Natural Telic, which we designate as T_N, namely the value *flow*. This is illustrated in (93) below.

(89) $$\begin{bmatrix} \textit{river} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{space} \\ \text{C} = \textbf{water} \\ \text{T}_N = \textbf{flow} \end{bmatrix} \end{bmatrix}$$

The naturally occurring activities described above are to be distinguished from those associated with intentionally created or designed objects, such as letters, pens, knives, etc.

Inherent in this is an association between the Agentive and Telic of the object, i.e., the object is made for a purpose (Agentive-Telic pairing), as in the case of *letter* in 2.1.4. Natural kinds lack this association, as they do not encode an Agentive value.

There are, of course, purposes and functions that can be attributed to natural kinds, most notably the recognition that something can be consumed as food, or ingested as a beverage. Hence, although we classify apples as fruits, they are, more importantly, *edible fruits*. Hence, just as we can "recover" the implied Telic for an artifactual noun, such as *cake* in (92a), we can perform the same computation in (92b), given our classification of apples as foods.

(90) $\begin{bmatrix} \textit{apple} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{fruit} \\ \text{T}_N = \textbf{eat} \end{bmatrix} \end{bmatrix}$

(91) a. Mary enjoyed the cake. (TELIC = eat)
b. Mary enjoyed the apple. (TELIC = eat)

In cases such as (92b), there is no Agentive-Telic pairing but there is an intention associated with the activity that has been given to the natural entity, namely, eating for sustenance and drinking for quenching, respectively, in the above examples. We will examine these cases in more detail in Chapter 5.

2.2.4 Agentive

In this section we examine in more detail the Agentive Quale, A, a role that encodes information about the origin of an object, or its "mode of coming into being", to use Aristotle's terminology. This is a crucial role for differentiating the kinds of objects, properties, and relations that exist in the world, since it provides a mechanism for distinguishing natural entities from non-natural entities. As cognitive agents interacting with the world, we can immediately distinguish between those objects that present themselves to us (occurring naturally) from the various artifacts that we create through our own activities and intentional behavior. This is captured in the qualia structure by reference to the Agentive role, whose default value of *nil* captures the primacy of a natural origin. Hence, the natural kinds *water*, *tiger*, and so forth, will have the following Agentive value:[4]

(92) $\begin{bmatrix} \textit{water} \\ \text{QUALIA} = \begin{bmatrix} \text{F/C} = \textbf{liquid} \\ \text{A} = \textbf{nil} \end{bmatrix} \end{bmatrix}$

By convention, when the value of a qualia role is *nil*, it will not appear as part of the qualia structure.

[4]It may be the case that specific aspects of how a natural object is brought into being can be represented in the Agentive, but we will not pursue this here. See also footnote 5.

Now consider how we can exploit this role to distinguish the different artifactual objects we encountered in the previous section, according to how they are created. Recall that nouns such as *cake* and *bread* were representing as denoting specific types of food, where the concept of *food* itself denotes a physical object with a Telic value making reference to the activity of eating. We can now make explicit reference to how an artifact is made by assigning that value to the Agentive, as illustrated in (94) for the noun *bread*.

(93) $$\begin{bmatrix} \textit{bread} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{food} \\ \text{T} = \textbf{eat} \\ \text{A} = \textbf{bake} \end{bmatrix} \end{bmatrix}$$

For many artifacts, the Agentive value is recognized or identified well enough to be paraphrased with a single activity. For example,

(94) a. $[house \quad [\text{A}= \textbf{build}]]$
b. $[painting \quad [\text{A}= \textbf{paint}]]$
c. $[letter \quad [\text{A}= \textbf{write}]]$
d. $[beer \quad [\text{A}= \textbf{brew}]]$

For some nouns, however, we may be able to identify (classify) them as artifacts without knowing their exact provenance. This is most likely the case for most complex artifactual objects we encounter in our daily lives, such those denoted by the nouns it computer, *car*, *refrigerator*, and so forth. In such cases, we identify to the best of our ability what the Agentive value should be. In these cases, a conventional speaker's lexicon would assign the underspecified activity of *make*, as shown below for *car*.

(95) $$\begin{bmatrix} \textit{car} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{vehicle} \\ \text{C} = \{\textbf{engine, door, wheels, ...}\} \\ \text{A} = \textbf{make} \end{bmatrix} \end{bmatrix}$$

The empirical evidence suggesting that certain nouns encode information about the origin of the object is somewhat similar to the one we examined for the Telic in 2.2.3. For example, the contexts in (114) clearly evoke a creation act made available by *thesis*, *sentence*, *movie* and *painting*, even if this act is not expressed syntactically. That is, the sentences in (114) are interpreted as referring to specific events associated with the direct object, namely the events expressing how this object came into existence.

(96) a. Paul completed his PhD thesis in 2000. (AGENTIVE = write)
b. She wouldn't let me finish my sentence.(AGENTIVE = speak)
c. Woody Allen has started a new movie. (AGENTIVE = direct, film)
d. John began a large oil painting yesterday. (AGENTIVE = paint)

It is interesting to note that the verb *complete* tends to select the Agentive value of its complements (114), while the verb *finish* may select either the Agentive or the Telic value, depending on the context. For example, in the VP coordination in (98a), the presence of the verb *publish* seems to force an interpretation of the Agentive quale (*write*), while in (98b), the possessive *his* in the object NP appears to block an Agentive reading, hence allowing the selection of a Telic interpretation.

(97) a. He just finished and published his first novel. (AGENTIVE = write)
b. I have just finished his first novel. (TELIC = read)

As mentioned above, nouns with an Agentive value denote artifacts, objects made by humans for a particular purposes. This may be a concrete physical object, as in the examples above, or an abstract object, as in the case of the noun *idea*, which refers to a proposition (*prop*) that was brought about by an act of thought.

(98) $$\begin{bmatrix} \textit{idea} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{prop} \\ \text{A} = \textbf{think} \end{bmatrix} \end{bmatrix}$$

However, in GL, also action nominals like *arrival* and *building* and agentive nominals like *violinist* or *singer* have an Agentive value. We will return to agentive nominals below, while we discuss action nominals in the section on Nominalization in Chapter 12.

Objects with an identical Formal Qualia value may have different Agentive values, if they differ in the way they came into being and if this difference is encoded in the lexicon for those concepts. For example, consider liquids such as *water* and *coffee*, which share a Telic value reflecting that they are potable liquids, as illustrated in (100) below.

(99) $$\begin{bmatrix} \textit{water/coffee} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{liquid} \\ \text{T} = \textbf{drink} \end{bmatrix} \end{bmatrix}$$

It is clear that the general function of these objects, i.e., their Telic role, does not distinguish them conceptually. This is accomplished by an additional dimension of classification, that of the Agentive quale, which specifies the mode of origin or creation (cf. (101)).

(100) $\left[\begin{array}{l} \textit{coffee} \\ \textsc{qualia} = \left[\begin{array}{l} \text{F} = \textbf{liquid} \\ \text{T} = \textbf{drink} \\ \text{A} = \textbf{brew} \end{array}\right] \end{array}\right]$

While coffee is a liquid that is created typically by the activity of "brewing", water is a natural kind which has been associated with a particular function. The identification of the Agentive role of "brewing" with the noun *coffee* acts to differentiate it from naturally occurring liquids. Notice how this is reflected in the interpretation of the Adjective-Noun construction in (102a) with *coffee*, as contrasted with that in (102b) with *water*.[5]

(101) a. fresh coffee (AGENTIVE = brew)
b. fresh water (in contrast to "salt water")

Similarly, the interpretation of the predicate *make* in the context of *cake* is 'bake', while in the context of *dress* it is 'sew'. This suggests that the Agentive values of *cake* and *dress* are $bake$ and sew, respectively.

(102) a. Mary made a cake. (AGENTIVE = bake)
b. Her mother made her a dress. (AGENTIVE = sew)

The interpretation of the Agentive Quale differs from the interpretation of the Telic in one important respect. While the event in the Telic has a *generic* interpretation (as we saw in 2.2.3) in the sense that it refers to a typical situation, such an interpretation is not available for the Agentive quale. In fact, the occurrence of the event specified in the Agentive is a precondition for the existence of the entity. Technically, we say that the Agentive event is an *existentially bound* or *existentially quantified* event, that precedes the existence of the object. On this view, one of the defining properties for the Agentive quale is that it presupposes the occurrence of the event it encodes.

An exception to this would be the analysis proposed in Busa (1996) for "role-defining" agentive nominals like *violinist*. According to her analysis,

[5]There are natural kinds which do permit modification by the adjective *fresh*, such as "fresh milk". Such cases involve an Agentive value that makes reference to activities that *enable* the object rather than *create* it. In this case, the activity of *milking* an animal. We will discuss these in Chapter 4.

a noun like *violinist* may be represented as denoting a person (F) having the ability (A) to play the violin (T). Under this analysis, the Agentive value of *violinist* encodes the precondition for the event specified in the Telic (i.e., the modal ability to play). The qualia structure representation proposed for the noun *violinist* is given below.

(103) $$\begin{bmatrix} \textit{violinist} & \\ \text{QUALIA} = & \begin{bmatrix} \text{F} = \textbf{human} \\ \text{T} = \phi\ \textbf{=play(human, violin)} \\ \text{A} = \textbf{ability(}\phi\textbf{)} \end{bmatrix} \end{bmatrix}$$

Besides playing a role as an implicit predicate, there are a number of constructions in which the Agentive event associated with a noun and the noun itself co-occur in context. Again, there is an important distinction to be made between constructions in which the event being expressed corresponds to the lexically specified Agentive value of the noun and those where the event being expressed updates the Agentive information of the noun in composition.

Verb-argument constructions in which the verb expresses the Agentive value of the noun typically consist of a creation predicate and its direct object, such as those illustrated below in (105).

(104) a. John baked a cake.
b. They built a house in Greece.
c. Mary took a photograph of her son.
d. Marc is painting a picture.
e. She wrote a letter to John.

An example of verb-argument construction in which the verb updates the Agentive value of the noun in composition is found in (106) where the predicate *type* expresses a manner of *writing*, that is, one of the possible ways in which the object *the letter*, is brought about, i.e., its Agentive value.

(105) She sat down and typed the letter.

It is interesting to note that Adjective-Noun constructions are ungrammatical when the adjective expresses the Agentive value of the nominal head. For example, consider the semantically anomalous pairs in (107) below.

(106) a. *baked bread (AGENTIVE = bake)
b. *a build house (AGENTIVE = build)
c. *a written book (AGENTIVE = write)

GL explains the anomaly of the examples in (107) as resulting from a lack of "informativeness" associated with the modification. Note that the expressions in (108) are anomalous because they are redundant:

(107) a. *a male bachelor
b. *a female woman

Similarly, modification by adjectives denoting the Agentive value of the head are also uninformative. If, however, additional information is given besides reference to the Agentive activity alone, the resulting expression is well-formed, as illustrated in (109).

(108) a. freshly baked bread
b. a well built house
c. a beautifully written book

Following this analysis, any semantic anomaly resulting from adjectival modification in Adjective-Noun constructions can be used as a diagnostic for identifying lexically specified Agentive values for nouns. Consider, for example, the expressions in (110):

(109) a. Sue wore a knitted sweater to the party.
b. His later works include painted portraits.

While the modifying adjectives *knitted* and *painted* in (110) express how the objects named by the head nouns were created, the resulting expressions are semantically acceptable and do not appear uninformative. From this evidence, we may conclude that the adjectives in (110) do not identify the lexically specified Agentive values of *sweater* and *portrait*, respectively. That is, by identifying the mode of creation for the sweater as knitting, the expression " knitted sweater" is informative, since there are several ways in which a sweater may be made. Similarly, a portrait may be created by the act of painting, photography, or perhaps other means. The expression "painted portrait" is hence informative.

2.3 Conventionalized Attributes

As we have seen throughout this chapter, Qualia structure is intended to provide a systematic and linguistically grounded representation of aspects of word meaning that are usually not modeled in linguistic theory. There are many properties and events, however, that are conventionally associated with an entity, but are not strictly part of the identified Qualia roles. In this section, we examine such information, which we will refer to as *conventionalized attributes* (CAs).[6] A conventionalized attribute is a property typically associated with an object through our *experiencing*, by means of various perceptual modalities, rather than of our *use* of it (but see the observations in (120) below). Examples of CAs can be seen in (111):

(110) a. Dogs bark.
b. The sun warms the air.
c. Water flows.
d. Airplanes make noise.

The notion of recording the properties conventionally associated with an object is, of course, related to efforts of encoding our commonsense knowledge of things in the world (Lenat, 1989, Hobbs et al, 1985, Havasi et al, 200X). The motivation behind the notion of a CA is to provide a lexical means for encoding default information that can be used in semantic composition within the sentence, as well as for logical inferencing above the sentence.

CAs may be activated in specific contexts in a similar fashion to Qualia. For example, the contexts in (112) clearly evoke a *sound* event made available by the referent of the nouns *dog*, *bird*, and *rain*, while those in (113) require that the referents of the nouns *flower*, *gas*, and *coffee*, have the ability of creating a perceivable *smell*. Within GL, it is assumed that both these activities are encoded in the semantics of the nouns in the form of CAs.

(111) a. They heard the village dog in the distance. (SOUND = barking)
b. Ann was listening to the birds. (SOUND = singing)
c. He could hear the rain in the garden. (SOUND = falling)

(112) a. John can smell the flowers in his garden. (SMELL = scent)
b. The repairman smelled gas in the kitchen. (SMELL = odor)
c. Mary woke up and smelled coffee. (SMELL = aroma)

Commonsense knowledge, such as that in (111a), provides the default values for contextualized interpretations, as in (112a).

It is a well-studied phenomenon that Qualia roles are accessed in certain contexts, such as those discussed earlier in the chapter, involving the

[6]See Pustejovsky, 2006, Pustejovsky and Jezek 2008, Pustejovsky and Rumshisky 2010 for first mention and development of this notion.

aspectual predicates *begin*, *start*, and *finish*. For example, the sentences in (114), repeated below, demonstrate how the Agentive role helps complete the interpretation.

(113) a. Paul completed his PhD thesis in 2000. (AGENTIVE = write)
b. She wouldn't let me finish my sentence. (AGENTIVE = speak)
c. Woody Allen has started a new movie. (AGENTIVE = direct, film)
d. John began a large oil painting yesterday. (AGENTIVE = paint)

But how can we account for the interpretations of the NP objects to the verb *enjoy* in (115)?

(114) a. Mary sat out and enjoyed the sun. (warming up)
b. It's a great place to enjoy the sea. (viewing, swimming, walking)

As with the aspectual predicates in (114), with the sentences in (115), the verb *enjoy* appears to reference an activity associated with the object, but not one that could be typically defined as filling the value of a Qualia role. Note that (115b) may invoke several activities made available by the noun's referent (*sea*-related activities). For example, as suggested above, it may evoke the activity of "viewing the sea", "swimming in the sea", "walking on the beach", and so forth. Within GL, it is assumed that the activities associated with *sea* may be classified according to a cline of conventionality, and that only the most conventional ones are coded in the noun's meaning as CAs. The details of the methodology adopted in GL to distinguish between coded vs. non-coded conventionalized activities will be presented in Ch. X. As we will see, this methodology is grounded on empirical evidence and focuses on distributional behavior as well as the comparison between elliptical and non-elliptical uses (for example, between "enjoy the sea" and "enjoy the sea view").

To see how CA interpretation differs from standard Qualia role values, consider the contextual interpretations we can associate with an adjective such as *fast*. When modifying artifactual nouns, it has long been observed that *fast* functions as an adverbial over the activity associated with the Telic role of the head noun it modifies.

(115) a. Mary is a fast typist. (TELIC = type)
b. This Porsche is a fast car. (TELIC = drive)

When there is no Telic role, however, the CA value(s) associated with the noun help provide an interpretation. Consider the sentences in (117) below.

(116) a. The tuna is one of the fastest fish in the sea. (swimming)
b. John was the fastest boy in the school. (running)

In GL it is assumed that conventionalized attributes are not external to Qualia Structure. Instead, they are interpreted as further characterizations of Qualia roles. That is, CAs are not seen as independent roles, but as projective manifestations of specific Qualia, which, together with the information coming from the Qualia, contribute to defining the overall semantic profile of a given type. The representation of CAs in GL is therefore always mediated through a Quale, as in (118) and (119) below.

(117) $$\left[\begin{array}{l} \textit{dog} \\ \text{QUALIA} = \left[\text{F} = \left[\begin{array}{l} \textbf{animal} \\ \text{CA} = \textbf{bark} \end{array}\right] \right] \end{array}\right]$$

(118) $$\left[\begin{array}{l} \textit{fish} \\ \text{QUALIA} = \left[\text{F} = \left[\begin{array}{l} \textbf{animal} \\ \text{CA} = \{\textbf{live\_in(water), swim}\} \end{array}\right] \right] \end{array}\right]$$

That is, the commonsense properties of "barking" relating to dogs and "swimming" relating to fish, is encoded within the Formal role, but identified as specific CA.

Up to this point we have focused on CAs as associated with nouns denoting natural kinds, but CAs can associate with artifacts as well. For example the sentences in (120) clearly involve reference to a sound made by the objects denoted by the nouns *car* and *airplane*, and the most plausible interpretations for these sounds are indicated in parentheses.

(119) a. I could hear a car behind me. (driving)
b. We do occasionally hear an airplane. (flying)

The GL analysis of the context in (120a) may be viewed as follows: we experience the *sound* of cars through their use; hence, the noun's Telic value *drive* interacts with the predicate *hear* to arrive at the conventionalized attribute for the car, i.e., the sound of it driving. The resulting representation is given below:

(120) $$\left[\begin{array}{l} \textit{car} \\ \text{QUALIA} = \left[\begin{array}{l} \text{F} = \textbf{vehicle} \\ \text{T} = \left[\begin{array}{l} \textbf{drive} \\ \text{CA} = \textbf{make\_noise} \end{array}\right] \end{array}\right] \end{array}\right]$$

2.4 Recursive Qualia Structures

Some scholars (Busa, 1996, Busa and Johnston, 1996, Bouillon, 1997, Pustejovsky, 1998) have suggested analyzing the individual Qualia roles as recursively making reference to qualia values themselves. One motivation for this comes from the interpretation of specific linguistic constructions including nominal compounds, where such recursive structure help capture the composition interpretations. For example, in the English compound *dining room* and the French compound *verre á vin* (wine glass), the Telic role for the entire expression is directly associated with "dining" and "wine" , respectively.

Now consider the following pair. The Telic role of the noun *rifle* is shown below as "firing":

(121) $\begin{bmatrix} \textit{rifle} \\ \text{QUALIA} = \begin{bmatrix} \text{T} = \textbf{fire} \end{bmatrix} \end{bmatrix}$

Notice, however, that in a compound construction such as *hunting rifle*, the Telic value for the compound is no longer that of "firing" but is the more specialized activity of "hunting".

(122) $\begin{bmatrix} \textit{hunting rifle} \\ \text{QUALIA} = \begin{bmatrix} \text{T} = \begin{bmatrix} \text{T} = \textbf{hunt} \\ \text{A} = \textbf{fire} \end{bmatrix} \end{bmatrix} \end{bmatrix}$

In such cases, the Telic is seen as having a more complex structure; namely, by firing the rifle (the local Agentive value), it can be used for hunting (the local Telic value).

2.5 Qualia Structure for Verbs

Our discussion so far has focused on Qualia roles as they relate to noun meaning. Qualia structure may be extended, however, as a representational mechanism for describing all the major linguistic categories, including verbs and adjectives. It is easy to informally identify the Telic of a sandwich (*eating*), the Agentive of a cake (*baking*), or the Constitutive of bread dough (*flour*); it might seem less obvious, however, what the corresponding Qualia are, for verbs denoting such different situations such as *building*, *walking*, and *breathing*.

For the domain of individuals, the Formal Quale acts to identify the subdomain within which an entity is placed. That is, we know beer is a liquid, rocks are physical objects, jokes are abstract objects, and so on. Nominals refer to stable objects, and as such can be organized hierarchically with persistent concepts, denoted by the Formal Quale.

For the concepts denoted by verbs, however, things are more complicated. Consider again the questions addressed earlier in the chapter introducing the notion of the Qualia roles.

(123) a. *Formal*: What kind of thing is it, what is its nature?

b. *Constitutive*: What is it made of, what are its constituents?

c. *Telic*: What is it for, how does it function?

d. *Agentive*: How did it come into being, what brought it about?

Modifying these questions to the domain of verb denotations, we see immediately a correlation between distinct Qualia roles and the different eventualities denoted by verbs. For example, the Formal can be seen as characterizing predicates denoting stable and persistent verbal predicates, namely states such as *love* and *believe*. But since verbs can denote changes in the world, they can refer to the manner in which something happens or changes, that is, the Agentive Quale. For example, intentional activities such as those denoted by the verbs *run* and *walk* can be characterized as Agentive Quale verbs. Change-of-state verbs such as *break* and *open* can be modeled as denoting a static resulting state (Formal) brought about by an activity (Agentive). On the other hand, intentional or directed events such as *build* and *clean* can be viewed as denoting a static intended goal state (Telic) brought about by an activity (Agentive). Hence, we can briefly identify verbs through their Qualia structure as illustrated below with specific examples:

(124) STATE:

$$\begin{bmatrix} \textit{love} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{love\_state} \end{bmatrix} \end{bmatrix}$$

(125) ACTIVITY:

$$\begin{bmatrix} \textit{run} \\ \text{QUALIA} = \begin{bmatrix} \text{A} = \textbf{run\_act} \end{bmatrix} \end{bmatrix}$$

(126) CHANGE_STATE:

$$\begin{bmatrix} \textit{break} \\ \text{QUALIA} = \begin{bmatrix} \text{F} = \textbf{broken} \\ \text{A} = \textbf{break\_act} \end{bmatrix} \end{bmatrix}$$

(127) ACCOMPLISHMENT:

$$\begin{bmatrix} \textit{build} \\ \text{QUALIA} = \begin{bmatrix} \text{T} = \textbf{build\_goal} \\ \text{A} = \textbf{build\_act} \end{bmatrix} \end{bmatrix}$$

In Chapter 4, we develop this model more fully, and show how Qualia structure interacts with event structure and argument realization.

2.6 Conclusion

Qualia Structure has predictive power with respect to the degree of conventionalization of certain activities. That is, activities that are related to the origin and purpose of an object are the most likely to be conventionalized and therefore included in a noun's lexical meaning.

Further Readings

Lexical representations based entirely or in part on qualia structure have been discussed and adopted in several other theoretical frameworks, including Jackendoff (1997, 2002), van Valin & Lapolla (1997), van Valin (2005, 2011), Cruse 2004 and Croft and Cruse 2004. From a computational perspective, Qualia have central in the work of Copestake (1993), Copestake and Briscoe (1991,1995), Bos, Mineur, Buitelaar (1993), Jensen and Vikner (1995), and the lexical database design and ontological architecture of the SIMPLE framework (Calzolari et al, 2000, Busa et al, 2001) and the Brandeis Semantic Ontology (Pustejovsky et al, 2006, Havasi et al, 2006).

Chapter 2

Semantic Coercion in Language: Beyond Distributional Analysis. *James Pustejovsky and Elisabetta Jezek*

Semantic Coercion in Language: Beyond Distributional Analysis

James Pustejovsky & Elisabetta Jezek

The distributional properties extracted from linguistic corpora for a word are regarded by many as the principle contribution to its meaning. While largely sympathetic to this view, we argue that lexical representations which are built from evidence of distributional behavior alone are unable to fully explain the rich variation in linguistic meaning in language. Lexical meaning is modulated in context and contextual semantic operations have an impact on the behavior that words exhibit: this is why a context-sensitive lexical architecture is needed in addition to empirical analysis to make sense of corpus data. As a case study that shows how distributional analysis and theoretical modeling can interact, we present a corpus investigation aimed at identifying mechanisms of semantic coercion in predicate-argument constructions, conducted within the Generative Lexicon (GL) model. GL theory is particularly suitable for this task, because it focuses on the many operations contributing to sentence meaning while accounting for contextual modulations compositionally. The analysis demonstrates the ubiquity of the phenomenon and highlights the limits of a theory-blind distributional analysis. In particular, it shows how coercion may alter the distributional behavior of words, allowing them to show up in contexts in which they would otherwise not appear. A descriptive theory of coercion as proposed here is relevant not only for theoretical considerations, but also for computational purposes such as the elaboration of annotation schemes for the automatic recognition and resolution of coercion phenomena in texts*.

1. Background and Motivation

There is a rich and growing literature of work in corpus-based and computational linguistics based on the distributional hypothesis in language (Harris 1954; for an overview, see Sahlgren 2006). A large body of work in language technology uses distributional information to compute semantic similarities between words. Various techniques are employed to translate distributional data into semantic representations and to clarify what kind of semantic knowledge is acquired through distributional evidence. Distributional evidence is currently used for a wide variety of tasks and applications, ranging from the construction of type systems, linguistic ontologies, computational lexical resources and so on.

Rivista di Linguistica 20.1 (2008), pp. 181-214 *(ricevuto nell'ottobre 2008)*

In the present work, we share the belief that contextual similarity relates to semantic similarity in some interesting way [1], and accept that the investigation of word distributional behavior constitutes an empirically well founded procedure to discover aspects of word meaning. Distributional analysis, however, underestimates the fact that a word's semantics may undergo modulations in composition, and that these modulations are not given *a priori* but depend on the contexts in which the word appears. This is why we claim that a distributional approach to word meaning representation is not sufficient. Since meaning is constructed compositionally, a lexical semantic model is needed to account for the fact that word behavior is not exclusively driven by inherent semantic properties but also adjusted by semantic compositional rules. In other words, lexical meaning is manipulated contextually and this problem cannot be ignored within a distributional approach to meaning acquisition and representation.

In this paper, we concentrate on the phenomenon of semantic coercion in predicate-argument constructions. We use coercion as a case study to show how distributional analysis is not able to fully capture the complexity of the semantic processes that take place in text, and why it cannot account for the mismatches between predicate and argument types that can be observed in corpus data. Also, we show that a lexical architecture such as GL is able to account for these problematic cases, since it embodies a dynamic representation of lexical meaning and foresees compositional rules which allow for type adjustments in context.

2. Theoretical Framework

Generative Lexicon (henceforth GL) aims to provide a compositional semantics for language that accounts for the contextual modulations in meaning that occur in real linguistic usage. That is, it can be seen as focusing on the distributed nature of compositionality in natural language. One important aspect of this "context modulation" is systematic polysemy. Recently, there has emerged an appreciation of how complex this problem is (Nerlich 2003), as well as a new understanding of the parameters at play in the interpretation of polysemous expressions. Within GL, two factors have been identified as contributing to the interpretation of polysemous terms: the nature of the expression's lexical semantic representation; and mechanisms for exploiting this information in context compositionally. In recent work, this distinction has been identified with *inherent* versus *selec-*

tional polysemy (Pustejovsky 2008). Indeed, polysemy cannot truly be modeled without enriching the various compositional mechanisms available to the language. In particular, lexically driven operations of coercion and type selection provide for contextualized interpretations of expressions, which would otherwise not exhibit polysemy. This is in contrast with Cruse's (2000) view that it is not possible to maintain a distinction between semantic and pragmatic ambiguity. Cruse suggests that polysemy is best viewed as a continuous scale of sense modulation. The view within GL is generally that a strong distinction between pragmatic and semantic modes of interpretation should be maintained if we wish to model the complexity and provenance of the contributing factors in compositionality.

The notion of context enforcing a certain reading of a word, traditionally viewed as selecting for a particular word sense, is central both to lexicon design (the issue of breaking a word into word senses) and local composition of individual sense definitions. However, most lexical theories continue to reflect a static approach to dealing with this problem: the numbers of and distinctions between senses within an entry are typically frozen into a grammar's lexicon. This *sense enumerative* approach has inherent problems, and fails on several accounts, both in terms of what information is made available in a lexicon for driving the disambiguation process, and how a sense selection procedure makes use of this information (cf. Pustejovsky & Boguraev 1993 for discussion).

When confronted by the messiness of corpus data, however, it can be difficult to see where lexical structure stops and context begins, in their respective contributions made toward building an interpretation. In this section, we confront this issue. First, we review our theoretical assumptions, and then outline the data structures and mechanisms responsible for the contextual modulations we will encounter from corpus data.

Classic GL (Pustejovsky 1995) proposes that a lexical item has available to it the following computational resources:

(1) a. Lexical Typing Structure: giving an explicit type for a word positioned within a type system for the language;
b. Argument Structure: specifying the number and nature of the arguments to a predicate;
c. Event Structure: defining the event type of the expression and any subeventual structure it may have; with subevents;
d. Qualia Structure: a structural differentiation of the predicative force for a lexical item.

The GL model defines a language for making types, where qualia can be unified to create more complex concepts out of simple ones. Following Pustejovsky (2001, 2006), the ontology divides the domain of individuals into three levels of type structure:

(2) a. NATURAL TYPES: Natural kind concepts consisting of reference only to Formal and Constitutive qualia roles;
b. ARTIFACTUAL TYPES: Concepts making reference to Telic (purpose or function), or Agentive (origin).
c. COMPLEX TYPES: Concepts integrating reference to the relation between types from the other levels.

Most early representations of GL lexical representations are grounded in terms of typed feature structures (Copestake et al. 1993, Bouillon 1997). The feature representation shown below gives the basic template of argument and event variables, and the specification of the qualia structure.

$$\begin{bmatrix} \alpha \\ \text{ARGSTR} = \begin{bmatrix} \text{ARG1} = x \\ \dots \end{bmatrix} \\ \text{EVENTSTR} = \begin{bmatrix} \text{E1} = e_1 \\ \dots \end{bmatrix} \\ \text{QUALIA} = \begin{bmatrix} \text{CONST} = \textbf{what } x \textbf{ is made of} \\ \text{FORMAL} = \textbf{what } x \textbf{ is} \\ \text{TELIC} = \textbf{function of } x \\ \text{AGENTIVE} = \textbf{how } x \textbf{ came into being} \end{bmatrix} \end{bmatrix}$$

The first two classes in (2) are defined in terms of qualia. For example, a simple natural physical object (3), can be given a function (i.e., a Telic role), and transformed into an *artificial type*, as in (4).

(3) $$\begin{bmatrix} \textbf{physobj(x)} \\ \text{FORMAL} = \textbf{physform(x)} \end{bmatrix}$$

(4) $$\begin{bmatrix} \textbf{artifact\_obj(x)} \\ \text{FORMAL} = \textbf{physform(x)} \\ \text{TELIC} = \textbf{Pred(E,y,x)} \end{bmatrix}$$

Artifactual types (the "unified types" in Pustejovsky, 1995) behave differently from naturals, as they carry more information regarding their use and purpose. For example, the noun *sandwich* contains information of the "eating activity" as a constraint on its *Telic* value, due to its position in the type structure; that is, **eat(P,w,x)** denotes a process, **P**, between an individual **w** and the physical object **x**. It also reflects that it is an artifact of a "making activity".

$$(5) \quad \begin{bmatrix} \textbf{sandwich(x)} \\ \text{CONST} = \{\textbf{bread,...}\} \\ \text{FORMAL} = \textbf{physform(x)} \\ \text{TELIC} = \textbf{eat(P,w,x)} \\ \text{AGENTIVE} = \textbf{make\_activity(z,x)} \end{bmatrix}$$

Complex types are reifications of multiple types, bound by a coherent relation. They are obtained through a complex type-construction operation on Naturals and Artifactuals. For example, *book* is a complex type denoting both the informational context and the physical manifestation of that content. One of the key properties of complex types is that they allow co-predication. In co-predication, two distinct senses of a lexical item are simultaneously accessed, for instance by applying two apparent incompatible types of predicates to a single type of object (as in 'the book I'm reading weights one kilo', 'the speech was long but interesting', etc.).

As mentioned above, there are two grammatical innovations necessary for enriching the model of selection. The first is a richer lexical representation, presented above. The second is a stronger theory of selection. Here we make reference to three mechanisms at work in the selection of an argument by a predicative expression (Pustejovsky 2008). These are:

(6) a. PURE SELECTION (Type Matching): the type a function requires is directly satisfied by the argument;
  b. ACCOMMODATION: the type a function requires is inherited by the argument;
  c. TYPE COERCION: the type a function requires is imposed on the argument type. This is accomplished by either:
    i. *Exploitation*: taking a part of the argument's type to satisfy the function;
    ii. *Introduction*: wrapping the argument with the type required by the function.

Given this three-way distinction, we can now ask when polysemy arises in grammar. We will argue that the ability to assign more than one interpretation to a lexical or phrasal expression is a result of *type coercion*. Lexical items that are inherently complex in their meaning, what have been termed complex types (or dot objects), will assume the interpretation of whatever selectional context they appear in (even if multiple contexts are available: see section 5.1.1. for fuller discussion). This phenomenon will be referred to as *inherent polysemy*, as the potential for multiple interpretations is inherent to the object itself. Most other cases of polysemy we will analyze as *selectional* in nature.

Now let us examine more closely the types in our language and the mechanisms at work in argument selection. From the point of view of their internal structure, Natural types (e.g. *lion*, *rock*, *water*) are atomic [2]. Conversely, artifactual (or tensor) types (e.g. *knife*, *beer*, *teacher*) have an asymmetric internal structure consisting of a *head type* that defines the nature of the entity and a *tail* that defines the various generic explanatory causes of the entity of the head type. Head and tail are unified by a type constructor ⊗ ("tensor") which introduces qualia relations to the head type: so, for instance *beer* = $liquid \otimes_{Telic}$ *drink*. Finally, complex types (or dot objects) (e.g. *school*, *book*, *lunch* etc.) have a symmetric internal structure consisting of two types clustered together by the type construction • ("dot"), which reifies the two elements into a new type. Dot objects are to be interpreted as objects with a complex type, not as complex objects. The constituents of a dot type pick up specific, distinct, even incompatible aspects of the object (for instance *lunch* picks up *event•food*, *speech* picks up *event•info* etc.) (more on this in section 4: as a general reference for the type syntax in GL, see Asher & Pustejovsky 2006).

The selection mechanisms introduced in (6) allow for modulation of types during semantic composition. *Matching* or *Pure Selection* takes place when the type call of the verb is directly satisfied by the argument. In this case, no type adjustment occurs. *Accommodation* occurs when the selecting type is inherited through the type of the argument. *Coercion* takes place when there is a mismatch (type clash) between the type selected by the verb and the type of the argument. This clash may fail to generate an interpretation (as in the case of '*The rock died'): if the verb is non-coercive, and the argument fails to pass the pretest imposed by the verbs type, it will not be interpreted by the interpretation function (the so-called *fail early* selection strategy – see Pustejovsky 2006). Alternatively, the type clash may trigger two kinds of coercion operations, through which the type required by the function is imposed on the argument type. In the first case, i.e. *exploitation*, a subcomponent of the argument's type is accessed and exploited (for example, in 'the author will discuss her book', *discuss* exploits the informational content of *book*), whereas in the second case, i.e. *introduction*, the selecting type is richer than the argument type and this last is "wrapped" with the type required by the function (for example, in 'the passengers read the walls of the subway', *read* "wraps" the *walls* with an informational content).

The reason why two coercion operations are proposed instead of one is that the information accessed to in semantic composition can be differently embedded in a noun's semantics. In both cases, however, coercion is interpreted as a typing adjustment. But where should the type adjustments take place, what sort of adjustments should be made and how pervasive is coercion? These are questions we address in the following sections.

3. Lexical Sets and Data Clustering

In our work we investigate the selectional behavior of types in text with the aim of detecting coercion phenomena and highlighting the inability of distributional analysis to fully capture the complexity of semantic processes occurring between types in composition; for this purpose, we observe the combinatorial 'space' of both verbs and nouns belonging to different types, focusing on the apparent mismatches between selecting and selected types. We adopt the methodology taken in Rumshisky et al. (2007) (see also Pustejovsky et al. 2004): we start by choosing a verb that selects for a given type, α; we automatically extract from our corpus the set of nouns (*lexical set*) that typically co-occur with this verb in a specified grammatical relation (for our current purposes, we restrict our investigation to the relation of *object-of* and, to a lesser extent, *subject-of*, although we are aware that coercions may apply to other relations as well, including indirect object and prepositional phrase)[3]. We then cluster those nouns into types (α_1, α_2, …) and distinguish those nouns satisfying the verb's selectional requirement from those which do not. Next, we evaluate what typing adjustments can apply to the residue noun set, in order to account for the underlying type mismatches, and how they should be represented. This procedure is repeated for a number of predicates selecting different types. We also carry out our investigation taking noun types as a point of departure. In this case, we follow roughly the same procedure: we select a noun of a given type, we extract the lexical set of verbs it combines with, we compare source and target types, we isolate the mismatches, and, finally, we speculate about the semantic operations at play in composition.

4. Beyond distributional analysis

When confronted with real corpus data, one can see at once how complex the clustering procedure is and how corpus investigation can not be conducted successfully without an appropriate architecture of the lexicon as a base. First of all, lexical sets don't map neatly onto semantic types. Consider for instance the verb *ring*. Typically, a person 'rings a *human*' (=call by telephone), but there are other entities which can be rang successfully in this verb sense, such as *institution*s and *location*s:

(7) *ring* (Body: 'call by phone'; Arg: *human*)
Object

a. *human*: mother, doctor, Chris, friend, neighbour, director
b. *institution*: police, agency, club
c. *location*: flat, house; Moscow, Chicago, London, place

Ex. I rang *the house* a week later and talked to Mrs Gould
The following morning Thompson rang *the police*
McLeish had rung *his own flat* to collect messages
I said Chicago had told me to ring *London*.

Next, lexical sets are not homogeneous paradigmatic structures. Instead, they seem to have core and satellite members (see Pustejovsky and Rumshisky 2008 & Rumshisky this issue). Consider for instance verbs that typically describe actions we do with *document*s (e.g. *read, publish, send, translate*). Although from a conceptual point of view *document* is a well-defined type, its linguistic membership seems to vary when we move from verb to verb (see Hanks and Jezek 2008 for discussion):

(8) What is a *document*?

a. *read*: <u>book</u>, <u>newspaper</u>, bible, <u>article</u>, <u>letter</u>, <u>poem</u>, <u>novel</u>, <u>text</u>, page, passage, story, comics script, poetry, <u>report</u>, label, verse, manual
b. *publish*: <u>report</u>, <u>book</u>, <u>newspaper</u>, <u>article</u>, pamphlet, edition, booklet, result, poem, <u>document</u>, leaflet, newsletter, volume, treatise, catalogue, findings, guide, novel, handbook, list
c. *send*: <u>message</u>, <u>letter</u>, telegram, copy, postcard, cheque, parcel, fax, card, <u>document</u>, invoice, mail, memo, <u>report</u>
d. *translate*: bible, <u>text</u>, instructions, abstract, treatise, <u>book</u>, <u>document</u>, extract, <u>poem</u>, menu, term, <u>novel</u>, <u>message</u>, <u>letter</u>

Finally, a word that is part of a lexical set may be an isolated item, in the sense that it may not fit in any of the individuated types. This is the case for instance of *chest* (*bodypart*) as object argument of *listen* (selecting for *sound*) (for an overview of argument types of *listen*, see section 6):

(9) 'your doctor will listen to your chest'

Given these observations, we ask: how can the data above be accounted for in a distributional model of the lexicon? How does distributional analysis account for the differences in argument type observed for *ring* within the same verb sense [4]? How does it deal with the problem of "shimmering" sets illustrated in (8)? We regard the conventional distributional view of the corpus as unsatisfactory. We argue that one of the reasons why sets and types do not overlap is because covert semantic mechanisms are at play in composition. In this view, imperfect mappings between sets and types contain potential candidates for coercion operations, and usage-based paradigmatic clusters of words, although necessary, are not sufficient to predict the meaning in context of complex linguistic expressions. In the following sections we present our corpus investigation as seen through the GL model.

5. A Typology of Coercions

In our investigation, we take as our point of departure previous research on compositional mechanisms in semantics and discourse (cf. Asher & Pustejovsky 2000 and 2006, Pustejovsky 2006), where a set of semantic typing adjustments and rules are developed in order to account for the mismatches between selecting and selected type.

Here, we adopt a simplified version of their analysis and use their predictions to guide our corpus investigation. We take into account the following aspects: 1) with artifactual and dot types, operations can affect the whole type or just one of its components; 2) coercions can be domain-preserving (for example from *entity* to *entity*) or domain-shifting (from *entity* to *event*), and level-preserving (from *artifact* to *artifact*) or level-shifting (from *natural* to *artifact*) (see Pustejovsky 2006). In both cases, what matters is if the domain or level of the coerced argument remains within the general domain or level of interpretation, or if it is shifted. If we focus on domain-preserving shiftings within the domain *entity*, and take into account the distinction between Natural, Artifactual and Complex types, the following operations are predictable (see Pustejovksy 2006):

Table 1. Verb-Argument Composition.

	Verb selects		
Argument is	**Natural**	**Artifactual**	**Complex**
Natural	Sel/Acc	Qualia Intro	Dot Intro
Artifactual	Acc	Sel/Acc	Dot Intro
Complex	Dot Exploit	Dot Exploit	Sel/Acc

Taking Table 1 as the starting point of our analysis, in the following section we present and discuss various instances of coercion that we detected using the methodology sketched in 3, seen from the point of view of the GL model. We then arrange them according to which type is coerced (Complex, Artifactual, Natural) and which adjustment is made (Exploitation, Introduction).

5.1. Type Exploitation

As mentioned above, within GL it is assumed that there are four computational resources available to a lexical item: *Type Structure*, *Qualia Structure*, *Argument Structure*, and *Event Structure* (cf. Pustejovsky 1995). In principle, there can be four corresponding sorts of exploitation: *TS exploitation*, *QS exploitation*, *AS exploitation* and *ES exploitation*. We focus here on Type exploitation and Qualia exploitation. These two are closely related, since as we clarified in 2. in our model Qualia are key constituents of the Artifactual type. We leave it to further work to investigate how Argument Structure and Event Structure exploitation works. Type Exploitation consists of exploiting part of internal structure of a given type. Since Naturals are atomic types with no internal structure, in principle they cannot be exploited in semantic composition (but see 5.1.3. for further comments). Conversely, Dot types and Artifactual types have an internal structure and can be exploited.

5.1.1. Dot exploitation

When an expression is typed as a dot object, such as *book* (*phys•info*), *house*, (*phys•loc*), *speech* (*event•info*) and *exit* (*event•loc*), it is disambiguated in context by the selecting predicative phrase, an operation we refer to as Dot exploitation. From the point of view of its computational cost, Dot exploitation is an inexpensive operation (i.e. a light form of coercion). It consists of exploiting one aspect of the complexity of a dot type (i.e. its inherent polysemy) by way of predicating over that aspect only (a predication also called *Object elaboration*: see Asher and Pustejovsky 2006, 14 and Asher, forthcoming). Dot exploitation can be left or right, depending on which aspect of the dot object is exploited: since in principle we assume that dot objects are commutative, from the point of view of their *modus operandi* the two operations are similar (but see additional remarks in Asher and Pustejovsky 2006). Examples of dot exploitations with the nouns mentioned above in object position are given in (10-13)[5]:

(10) *book* (*phys•info*)
Object

a. *phys*: close, open, shut, throw away, steal, keep, burn, put away, bind, design, store, grab, drop, destroy, dust, hold, shelve, pile, store
b. *info*: ban, consult, edit, find interesting, study, translate, review, love, judge, revise, examine, like, describe, discuss

Ex. Jess almost *dropped* the book, then hastily replaced it on the shelf
The author *will be discussing* her new book

(11) *house* (*phys•location*)[6]
Object

a. *phys*: built, buy, sell, rent, own, demolish, renovate, burn down, erect, destroy, paint, inherit, repair
b. *location*: leave, enter, occupy, visit, inhabit, reach, approach, evacuate, inspect, abandon

Ex. They *built* these houses onto the back of the park
The bus has passed him as he *left* the house

(12) *speech* (*event•info*)
Object

a. *event*: deliver, make, give, finish, interrupt, conclude, end, begin, start, complete, cut (short), open
b. *info*: analyse, interpret, understand, quote, applaud, criticize, condemn, revise, translate, oppose, appreciate

Ex. He was forced to *interrupt* his speech while order was restored
US officials *condemned* the speech

(13) *exit* (*event•location*)
Object

a. *event*: make, facilitate, follow, force, hasten, register
b. *location*: block, bar, take, find, mark, indicate, reach, choose, locate

Ex. I very swiftly *made* my exit through the door
She *was blocking* the exit of a big supermarket

Examples (10-13) show that the single aspects (senses) of a dot object are often picked up separately. Many lexical items which are typed as dots tend to show up in text in just one of their aspects instead of both[7]. There are often asymmetries of use in dot exploitations, i.e. selectional preferences for one of the constituents (or aspects, or senses) of the complex type. Asymmetries may be within the same argument position, as noted by Jezek & Lenci (2007) with respect to the object position of the complex type *phys•info* (i.e. *letter*, *article*, *book*, *novel* etc.): It. *articolo* 'article' for instance combines more frequently with *info*-selectors rather than with *phys*-selectors:

(14) *articolo* 'article' (*phys•info*)
Object

a. *phys*: spostare 'move', ritagliare 'cut out'
b. *info*: approvare 'approve', bocciare 'reject', citare 'quote', correggere 'correct', ignorare 'ignore', commentare 'comment', conoscere 'know', condividere 'share'

Ex. *Ritaglia* tutti gli articoli che lo riguardano
'He cuts out all the articles about him'
Condivido interamente il suo articolo
'I agree entirely with his article'

Jezek & Lenci (2007) also note that lexical items realizing the same dot type exhibit interesting variations as far as their asymmetry goes: for example in object position *romanzo* 'novel' avoids the *phys* sense more than *libro* 'book' does.

(15) *romanzo* 'novel' (*phys•info*)
Object
a. *phys*: collocare 'place', portare 'carry'

(16) *libro* 'book' (*phys•info*)
Object

a. *phys*: bruciare 'burn', portare 'carry', distruggere 'destroy', rubare 'steal', conservare 'keep', custodire 'keep', buttare 'throw away' [8]

The same holds for *articolo* (fewer *phys* selectors) and *lettera* 'letter' (more *phys* selectors):

(17) *articolo* 'article' (*phys•info*)
Object

a. *phys*: spostare 'move', ritagliare 'cut out'

(18) *lettera* 'letter' (*phys•info*)
Object

a. *phys*: imbucare 'post', conservare 'keep', infilare 'put', distruggere 'destroy', raccogliere 'pick up', esibire 'exhibit, show', ritrovare 'find again', perdere 'lose', portare 'bring'

Ex. *Raccolse* la lettera da terra
'He picked up the letter from the ground'

Asymmetries of use may also be related to specific argument positions. With respect to the (pseudo-)dot type *animal•food* (i.e. *chicken*, *lamb* etc.) Rumshisky et al. 2007 note for instance that the subject position tends to disprefer the *food* sense, whereas this same sense dominates in the object position. A similar asymmetric behavior is found with *producer•product*, where the subject position tends to not select the *product* sense:

(19) *Honda* (*producer•product*)
Subject

a. *producer*: design, build, produce, create, assemble, accept, invest, work on, hate, introduce, develop, win, support, announce, invest, declare, say, acquire, be confident, be grateful, withdraw, bring out, decide, run, threaten, sponsor
b. *product*: stand, spin out of control, go on sale, be a missile

Ex. Honda immediately *withdrew* the two affected models
Their Honda *spun out of control*

Asymmetry of use can be a generic property of some dots, no matter what argument position they occupy. Both *door* and *gate* (*phys•aperture*) show preference for the *phys* interpretation in all argument positions (we restrict our example to *door*):

(20) *door* (*phys•aperture*)
Object

a. *phys*: slam, push, pull, bang, kick, knock at, smash, hold, paint, hit, remove, damage, replace, decorate
b. *aperture*: pass, enter, block

Subject

a. *phys*: swing, bang, shake
b. *aperture*: lead, go, give access, connect

Ex. Somewhere in the house a door *slammed*
The main door *went into* a small lobby

Interview (*event•info*) shows a distinct preference for the *event* interpretation in both subject and object position:

(21) *interview* (*event•info*)
Object

a. *event*: conduct, give, arrange, attend, carry out, terminate, conclude, close, complete, end, hold, cancel, undertake, extend, control, continue, begin
b. *info*: structure, discuss, analyze, describe

Subject

a. *event*: last, go well, take place, follow, end, progress, begin, become tedious, precede, start, happen
b. *info*: covers, centre on, concern, focus on

Ex. Officials *will be conducting* interviews over the next few days
Let's *discuss* the interview

Asymmetries of use as found in the corpus may be seen as an additional diagnostic together with co-predication for identifying dot objects [9]. While co-predication motivates the existence of dot objects, the asymmetry of use questions their 'dottiness' and hints that they might be types with an asymmetric internal structure, i.e., Artifactual types. It is not clear, however, if that is the case (if asymmetries of use really question dottiness). Firstly, asymmetry of use reflects usage and although usage is a key indicator of linguistic organization, it is an indirect one. Secondly, non-lexical factors may be relevant, such as the well-attested preferential linking between subject position and semantic components like animacy (the *animal* sense of *chicken*) and volitionality (the *producer* sense of *Honda*).

5.1.2. Artifactual Exploitation

Instances where an artifactual type is exploited only partially in composition can be accounted for as operations of Artifactual (or Tensor) exploitation. If only the head of the type is exploited no true coercion occurs: the selecting type is inherited through the type of the argument and the operation amounts to a sort of type accommodation. This occurs for instance when a verb selecting for a natural type (*fall*, *die*, *flow*) combines with an artifactual entity and selects only for the head of the type (cf. Table 1) [10].

(23) a. The *pen* fell to the floor
b. The *roof* has fallen and should be replaced

Conversely, if only the tail of an artifactual type is exploited (Qualia Exploitation), a coercion occurs. The value of the Quale of the argument is lifted into the type structure and then exploited in semantic composition [11]. This occurs for instance when an aspectual verb like *finish* (which types its internal argument as *event*) combines with an artifactual entity. First, the verb introduces an *event* (Event Introduction, henceforth E-I); then, as a response to the type call of the verb, the value of the Qualia is lifted at the level of interpretation (Qualia Exploitation, henceforth Q-E). In this way, the artifactual entity can be coerced to the type *event* and successfully fill the verb's argument slot [12].

(24) *finish* (Body: 'bring to an end'; Arg: *event*)
Object

a. *event*: journey, tour, treatment, survey, race, game, training, ironing, shopping
b. E-I, Q-E of phys$\otimes_{telic}$ τ: penicillin, sandwich, cigarette, cake, dessert, food
c. E-I, Q-E of liquid$\otimes_{telic}$ τ: drink, wine, beer, whisky, coke

Ex. When they finished the *wine*, he stood up
Just finish the *penicillin* first

What is significant here is that the meaning of *finish* ('bring to an end') is quite similar, regardless of the semantic type of the internal argument it appears with: in all examples, the bringing to an end of an eventuality is at stake. The meaning in context, however, (the co-compositional interpretation of the verb with its argument) will allow modulations in meaning, depending on the semantics of the object [13].

Finish is a 'strong' coercive verb, i.e. many of its objects are not pure events but rather dots or artifactuals (we restrict our observations to artifactuals here) [14]. This is not a characteristic of aspectual verbs in general: some aspectual verbs just don't coerce their arguments or they do it to a lesser extent. *Last* exhibits a few artifacts as subjects, and they are all re-interpreted as the interval of time for which their function holds:

(25) *last* (Body: 'occur over a certain time span'; Arg: *event*)
Subject

a. *event*: marriage, trial, siege, honeymoon, war, journey, strike, storm, rainfall
b. E-I, Q-E of phys$\otimes_{telic}$ τ: battery, cartridge

Ex. The *battery* lasts 24 hours
The *cartridge* lasted three weeks

Many non-aspectual event selectors (such as *attend, avoid, prevent, cancel, delay, schedule, skip* etc.) are 'weak' coercive verbs (i.e. the vast majority of their arguments are *event*s: in principle, those which are not, are coerced - but see section 5.1.2 for further discussion):

(26) *attend* (Body: 'be present at'; Arg: *event*)
Object:

a. *event*: meeting, wedding, funeral, mass, game, ball, event, service, premiere
b. E-I, Q-E of location$\otimes_{telic}\tau$: clinic, hospital, school, church, chapel

Ex. About thirty-five close friends and relatives attended the *wedding*
For this investigation the patient must attend the *clinic* in the early morning
He no longer attends the *church*

Again, one might argue here that *attend* does not exhibit the same meaning in all these contexts, and that a new meaning is licenced when *attend* occurs in combination with *location*s ('go regularly to'). In contrast to this view, we claim that the meaning of *attend* is much the same in all examples in (26). Also, we argue that the nouns *clinic*, *school*, *church* etc. are all successfully coerced to *event* because they denote functional locations associated to specific activities coded as Telic values (medical treatment, class, mass and so on). It is to these activities that we refer to when we say that we attend such locations: in other words, the combination of *attend* with a functional location ends up meaning 'to be physically present at an event in a given location' [15]. A similar argument applies to *avoid*, where the physical object *food* for instance is re-interpreted as the event of eating it [16]:

(27) *avoid* (Body: 'keep away from, stop oneself from'; Arg: *event*)
Object:

a. *event*: collision, contamination, clash, damage, accident, pregnancy, injury, question, arrest, starvation, war
b. E-I-Q-I of phys$\otimes_{telic}\tau$: food
c. E-I-Q-I of abstr$\otimes_{telic}\tau$: tax
d. E-I-Q-I of location$\otimes_{telic}\tau$: prison

Ex. Try to avoid fried *food*
You can't avoid the inheritance *tax* in those circumstances
His wife avoided *prison* because she is five months pregnant

Similarly to aspectual verbs like *finish* and event selectors like *attend*, perception verbs like *hear* may exploit the Qualia values of their internal arguments, if those are entities whose primary function (purpose) is to emit a sound (*bell*, *siren*, *alarm clock* etc.):

(28) *hear* (Body: 'perceive with the ear'; Arg: *sound*)
Object

a. *sound*: voice, sound, murmur, bang, thud, whisper, whistle
b. Q-E of phys$\otimes_{\text{telic}}\,\tau$: siren, bell, alarm clock

Ex. Then from the house I heard the *bell*
You can hear *sirens* most of the time
The next thing he heard was his *alarm clock*

Additional data of exploitation of Qualia values of artifactual types, as found in the corpus, are shown below. In the examples, selectors are grouped together according to the coercion operation at play (indicated by E-I, Q-E), instead of their semantic type. When no indication is present, we assume the operation at play is type matching or pure selection as in (29b) or type accomodation as in (29a):

(29) *bell* (*phys*$\otimes_{telic}$ *ring*)
Subject

a. *phys*: hang, swing, weigh
b. *phys*$\otimes_{telic}$ *ring:* sound, tinkle, clang, echo
c. Q-E, where *telic* = *ring*: awaken, interrupt, alert, warn, disturb, announce
d. E-I, Q-E, where *telic* = *ring*: begin, stop, start

Ex. The bells *warned* the inhabitants of the villages
It was at just that moment the bells *began*
When the bell *stopped*, we all went into lines

(30) *sandwich* (*phys*$\otimes_{telic}$ *eat*)
Object

a. *phys*: grab, fold, wrap
b. *phys*$\otimes_{telic}$ *eat*: munch, devour, chew
c. E-I, Q-E, where *telic* = *eat*: finish, refuse, mind, abandon, enjoy, try, avoid

Ex. I'm sure David *won't mind* sandwiches for a day
I *abandoned* a perfectly good bacon sandwich

(31) *wine* (*liquid*$\otimes_{telic}$ *drink*)
Object

a. *liquid*: pour, spill
b. *liquid*$\otimes_{telic}$ *drink*: drink, sip, gulp (down), down
c. E-I, Q-E, where *telic* = *drink*: finish, enjoy, prefer, try

Ex. Clarissa nervously *sipped* her wine
They *had finished* the wine and talked about almost everything

(32) *glass* (*phys*$\otimes_{telic}$ *hold (liquid)*)
Object

a. *phys*: raise, clink, lift, break, put down, clean, hold, set down, throw
b. *phys*$\otimes_{telic}$ *hold(liquid)*: refill, fill, empty
c. AS-E: where *telic* = *hold*(*liquid*): drink, pour, down, swallow
d. E-I, Q-E, where *telic* = *hold*(*liquid*), AS-E: finish

Ex. As a rule he only drank one glass, but that night he *drank* three
She *poured* two glasses and gave him one
When she*'d finished* the second glass, he was still there [17]

As we can see from the examples above, Qualia exploitation is ultimately an operation which lifts semantic information coded in the Qualia at the level of interpretation, as a response to a call of the verb for the type *event*.

Not all artifactuals are Q-exploited, however. Some artifactuals enter coercive contexts less easily than others. For instance, functional objects like *knife*, *car*, *pen*, *bed*, *table*, as opposed to *food* products and *drink*s in (24), are not often coerced to the events they typically participate in (*cut*, *drive*, *write*, *sleep*, *support*, respectively). Still, this does not mean that they do not undergo other kinds of coercion operations: for example, the noun *table* (physical object) may be successfully coerced to *location* (cf. 5.2.). This suggests that there may be conditions on coercion of artifactual types to events [18]. Also, this suggests that generative rules like event type coercion may apply semi-productively in a fashion similar to processes of word formation, which are regular but not systematic in their application.

Like Dot Exploitation, Qualia Exploitation is an ampliative rule which preserves the type structure but triggers the addition of new information to logical form (cf. Asher & Pustejovsky 2006). However, Qualia Exploitation differs from Dot exploitation because the inference it permits can be overridden in context (i.e., a different inference can be imposed contextually – see Lascarides & Copestake 1998):

(33) 'I ought to cancel the milk tomorrow.'

Milk is a liquid to be drunk (what in Pustejovsky 2008 is called a "natural functional type"): we would expect coercion to exploit the drink activity specified in the Telic Quale value (as in 'finish the milk'): however, the predicate *cancel* overrides this value and introduces a different inference (the delivery).

Qualia exploitation is more 'internal' than Dot exploitation and computationally more expensive. The disambiguation between Dot Exploitation and Qualia Exploitation follows from the way we structure the type associated with the noun. The two options available (Dot or Artifactual type) differ exclusively in the way a specific piece of semantic information is encoded: either as a type subcomponent (for instance *bottle* (*container•containee*) or as Qualia value (*bottle container*$\otimes_{telic}$ *hold(liquid)*). In order to assign a type to a term, we analyze its combinatorial behavior looking at the ontological and semantic properties of the words it typically combines with [19].

Following Asher & Pustejovsky 2006, not only dot objects but also artifactuals allow co-predication, since the NP denotation is embedded within the coerced interpretation (see also Copestake & Briscoe 1995, p. 13).

(34) She *opened* the wine and *poured* some into the glass

In (34), two senses of *wine* (*liquid* and *container*) are activated simultaneously in context. However, we assume that while the *liquid* sense is inherent, the *container* sense is introduced contextually by the verb (*open*) [20]. Thus, while it appears both dot objects and artifactuals allow co-predication, this is possible only under coercive contexts for artifactuals, such as that seen above with *open*. This does, however, make the distinction between them harder to characterize distributionally in some cases.

5.1.3. Natural Exploitation

In this section, we explore briefly how aspects of Natural types are referenced in various selectional contexts. Since a Natural type is atomic, any type exploitation performed over it is formally identical to type matching (i.e., pure selection). Yet it is apparent that some Naturals carry information about their prototypical use (e.g., water is for drinking) while other naturals do not (e.g., rocks are not *for* anything specific). We assume, however, that such information in Naturals is encoded not as qualia, but rather is associated with specific qualia as *conventionalized attributes* (Pustejovsky, 2008). A conventionalized attribute (CA) is a property we associate with an object through our *experiencing* of it, through various perceptual modalities, and not necessarily our *use* of it. For example, it is a property of most animals that they produce specific sounds, and this attribute can be invoked by perception predicates like *hear* and *listen*, which select for the type *sound*:

(35) Ann *was listening* to the birds (singing)
They *heard* the village dog in the distance (barking)

Similarly, we have conventionalized values associated with natural force event nominals, such as *wind* and *rain*:

(36) He could *hear* the rain in the garden (falling)
I couldn't *hear* anything but the wind in the trees (blowing, howling, whistling)

Thus, conventionalized attributes are typical properties of entities and may play a role in composition processes. Their role in composition, however, is different from the role played by the Telic and Agentive Qualia. While the latter may act like tensors and shift a type from natural to artifactual, the former may not [21].

Conventionalized attributes may associate with Artifactual types (like *car*) and Dot objects (like *door*) as well: basically, it may apply to all kinds of objects, under the appropriate circumstances. For example:

(37) Alice had heard the *car* and came out to him from the kitchen
Alan heard a *door* a few minutes before he last looked at his watch

It should be noted that we experience the sound of cars through their use, so the noun's Telic value interacts with the selecting type from *hear* to arrive at the conventionalized attribute for the car, i.e., the sound of it driving [22].

Although we will not explore the specific mechanisms responsible for this composition here, it is important to note that conventionalized attributes constitute an additional resource available to a defeasible semantic interpretation, in addition to values from the qualia structure. It remains an empirical question whether such attributes should be considered information associated with a lexical item or as purely ontological properties which, if violated in composition, give rise to a conceptual conflict which fails to licence an interpretation (on conceptual conflicts and consistency criteria, cf. Prandi 2004). Also, it is not completely clear if conventionalized attributes are external to Qualia Structure or if they are part of it (for example, if they are a further characterization of the formal and/or the constitutive role). Whatever the case may be, they enrich the material with which compositional mechanisms may apply.

Although it is still somewhat unclear, the interpretation of conventionalized attributes is most likely a coercion operation; but it both *introduces* a type, *sound*, while also potentially exploiting a value associated with the head being coerced. Note, however, that while one can 'smell a rose' because it is a formal attribute of most flowers to emit a scent, one does not typically 'smell a table', because this attribute is not normally true of physical objects like tables. Observe below the selectional behavior of the verb *smell* as seen in the corpus (CA-I stands for Inheritance of Conventionalized Attribute):

(38) *smell* (Body:'perceive or detect by the faculty of smell'; Arg: *odour*, *scent*)
Object

a. *odour*, *scent*: scent, perfume, fragrance, smell, odour, aroma
b. CA-I: smoke, soap, flower, whisky, gas, coffee, sea, petrol, cooking, bacon, dog, rose, food, drink, wine

Ex. I took a deep breath and smelt the *sea*
Smell the wild *flowers*
I can smell *gas*! Can you?
Wake up and smell the *coffee*

In (38), we find arguments of different type levels (naturals, artifactuals) and different type sorts (liquids, food etc.), and these arguments are all interpreted as *scents* or *odours* when appearing as the object of *smell*, since this is the selecting type. The specific interpretations arrived at in these sentences are made possible by the fact that the property of emitting a scent is a conventionalized attribute of all these objects and can be interpreted in semantic composition as result of *scent* applied to that object denotation [23].

5.2. Type Introduction

Instances where conceptual material is introduced, which is not part of the original meaning of the word, can be accounted for as operations of Type Introduction or Qualia Introduction. In computational terms, Introduction is an expensive operation if compared with Exploitation. Instead of exploiting a subcomponent of the argument's type, Introduction "wraps" the type of the argument with the type required by the function and makes new conceptual material available to interpretation. We have already seen several examples of Event Introduction in section 4.1.2., triggered by aspectual verbs like *finish* and event selectors like *attend* when they combine with artifactual types such as *food* or *functional locations*. Additional examples of Introduction as found in the corpus are illustrated below:

(39) *open* (Body: 'cause to become open'; Arg: *container*)
Object

a. *container*: drawer, bottle, cupboard, envelope, folder, tin, can, box, fridge, bag, cage, suitcase
b. *liquid*: wine, champagne, beer

Ex. I opened the *wine* carefully
Just as he was about to open the *beer*, the doorbell rang

(40) *leave* (Body: 'go away from'; Arg: *location*)
Object

a. *location*: room, house, country, England, flat, island, pub, kitchen, shore, station
b. *event*: concert
c. $phys \otimes_{telic} \tau$: table, car

Ex. He left the *concert* early
He left the *table* without taking lunch
I just left my *car* and ran

As we said above, Introduction adds new information which is not part of the noun's original meaning (for instance, *container* is not part of the meaning of *wine* and *location* is not part of the inherent meaning of *concert*). However, not all introductions are possible; for an Introduction operation not to fail it is important that the new information is semantically compatible with the lexical representation of the object and with its ontological properties. In (39), for instance, Introduction is successful because *wine*, *beer* etc. are artifactual liquids typically stored in containers. In (40b-c) *location* is introduced successfully because a *concert* is an *event* which takes place in specific locations and a *table* is an artifact around which people gather and spend time for specific purposes. Finally, in (40c) the *location* where the *car* is parked is introduced contextually by the function (*leave*).

As we can see, the distinction between metaphysics and the lexicon is again very relevant. Does Introduction lift to logical form something which is coded in the lexicon or does it exploit our world knowledge about the coerced entity? How can we possibly draw a line between these two options and is it necessary? These questions are not so easily answered when real corpus data are encountered. Consider again the verb *ring* in (41) (adapted from 7):

(41) *ring* (Body: 'call by phone'; Arg: *human*)
Object

a. *human*: mother, doctor, Chris, friend, neighbour, director
b. *location*: flat, house; Moscow, Chicago, London

Ex. I rang the *house* a week later and talked to Mrs Gould
I said Chicago had told me to ring *London*.

In (41b) is the type *human* introduced or is it exploited? We believe that in this case the operation at play is exploitation, since *house*, *Chicago* and *London* denote functional locations where people live or work and this information is most likely coded in the Qualia values of these expressions.

5.2.1. Dot Introduction

Instances where a predicate selecting for a dot type combines with an argument which is a non-dot can be accounted for as operations of Dot Introduction. In this case, the predicate coerces the argument type to dot object status. Examples of dot introduction are provided by the verb *read* which selects a *phys•info* type as internal argument but exhibits also non-dots in object position:

(42) *read* (*phys•info*)
Objects

a. *phys•info*: book, bible, article, brochure, letter, note, novel, text, document, diary, manuscript, manual, telegram, mail, pamphlet, hand-out; label, meter, timetable, sign
b. *info*: list, news, inscription, sentence, content, writing

Ex. I've come to read the *meter*
He could just read the faded *inscription* painted above the window

Some of the object arguments of *read* fail to match any of the subcomponents of the *phys•info* type; in such cases, the whole complex type is imposed on the source type:

(43) *read* (*phys•info*)
Objects

a. *phys*$\otimes_{telic}$ *write*: Dante, Proust, Homer, Shakespeare, Freud

Ex. That is why I read *Dante* now

The predication 'read Dante' is felicitous because the type of the argument is *human* agent of writing activity. *Read* also exhibits arguments which are dot objects but match the required type only partially:

(44) *read* (*phys•info*)
Objects

a. *event•info*: story, description, judgement, quote, reply, speech, proclamation, statement, question, interview
b. *sound•info*: music

Ex. I've read your *speeches*
I discovered he couldn't read *music*

In the examples above, the subcomponent *phys* (absent in the noun type) is introduced contextually. In:

(45) I tend not to read *long interviews* with top celebs

read introduces the *phys* component (not inherent in the noun *interview* – which type is *event•info*), while *long* exploits the *event* type.

All the above arguments of *read* are coerced to *phys•info* status: however, since the source types of the nouns can differ, different sorts of introductions take place. Instead of being coerced to the *phys•info* type, some of the arguments of *read* may license a shift in the verb's meaning, resulting in a more extended or metaphorical sense, as shown below.

(46) 'decipher' sense:
I can't read your *handwriting*
The *code* can be read properly

(47) 'interpret' sense:
He read her *expression* correctly
I wish I could learn to read those early prophetic *signals*
He must have read my *thoughts*

5.2.2. Qualia Introduction

When a verb selecting for an artifactual type combines with a natural type and coerces it to a certain function or purpose, Qualia (or Tensor) Introduction occurs. *Eat* and *drink* provide examples of this:

(48) *eat* ($phys \otimes_{telic}$ *eat*)
Object

a. $phys \otimes_{telic}$ *eat*: sandwich, pancake, bread, biscuit, pie, cake, steak, toast, ice-cream, snack, pudding, salad, meat
b. *phys* (natural): fish, chicken, worm; apple, banana, orange; mushroom, lettuce, spinach; grass, leaf, hay; fat, nut, rice, flesh

Naturals co-occurring with *eat* in object position are entities of different types (animals, fruits etc.): in the context of *eat* all these entities are re-computed as edible objects.

(49) *drink* ($liquid \otimes_{telic}$ *drink*)
Object

a. $liquid \otimes_{telic}$ *drink*: beer, wine, champagne, juice, sherry, lemonade, coke
b. *liquid* (natural): blood

Ex. Fanatics have been drinking horses' *blood* to gain strength

Blood is a *liquid* but it is not meant to be drunk: it can however be re-interpreted as *beverage* (*liquid*$\otimes_{telic}$ *drink*) contextually. Qualia introduction endows a Natural entity with a specific use (purpose) and shifts its type from Natural to Artifactual (cf. Pustejovsky 2006).

Qualia Introduction differs from Qualia Exploitation because the inference it permits is not inherited lexically. By definition, naturals do not carry prior information to suggest what their interpretation may be in a coercive environment, and their interpretation is strictly dependent on a specific context. If we examine the naturals appearing as direct objects of *eat* and *drink*, however, we may note that some of them are more easily reinterpreted as food or beverage than others (compare *water* and *milk* vs. *blood*). As we already clarified in 5.3.2., this occurs because even if Naturals do not have a complex Qualia Structure as Artifactual types do, some of them may exhibit inherent conventional attributes and natural telic aspects which may be exploited in semantic composition.

6. The Scope of Coercion Operations

In the previous sections, we have analysed in detail various kinds of compositional mechanisms of argument selection as they emerge from corpus data. We have distinguished between two main sorts of coercion operations, i.e. Exploitation and Introduction. We have also observed that verbs may vary with respect to their coercion potential, and that some nouns enter coercive contexts more easily than others.

In what follows, we take a broader perspective on coercion phenomena: that is, we evaluate briefly what the 'span' of coercion mechanisms may be, i.e., what semantic or conceptual shifts are possible (given a certain starting point); what can be coerced into what else; how easily this may occur etc. We assume that this span can be 'measured' by comparing the type selected by a given predicate (*target type*) with the list of argument types it occurs with in texts (*source types*). From the point of view of cognitive and psychological studies, as well as linguistic theory, these are all very interesting questions. A cartography of coercions based on the comparison between source and target types would give us much insight into human conceptualization and its generative nature.

Let us consider, for instance, the verb *listen* and assume it selects for *sound*. Corpus data show that *listen* combines with an extremely wide variety of arguments, only a subpart of which are *sound*s or *sound*-related types, i.e. types in which the sound dimension is coded lexically as a constraint to a Qualia value or as a conventionalized attribute (in 50 we restrict ourselves to a selection of these types):

(50) *listen* (*sound*)
Object

a. *sound*: voice, noise, ticking, hum, echo, hiss, thud, roar
b. *sound•info*: music, jazz; concert, opera, overture, tune, lyric, song
c. *event* (natural): rain, wind
d. *event* (involving sound production): breathing, whisper, cry; footstep
e. *event•info* (speech act): announcement, conversation, discussion, debate, speech, talk, dialogue
f. $phys \otimes_{telic} play$ (*sound•info*): radio, stereo
g. *phys•music*: disc, tape, record, album, cassette
h. $phys \otimes_{telic} ring$: bell, clock
i. $human \otimes_{telic} sing$, $human \otimes_{telic} speak$: singer, speaker
j. $human \otimes_{agent} write$ (music): Beatles, Mozart, Wagner, Bach
k. *human*: colleague, nurse, costumer, parent, friend
l. *phys* (body part): chest, heart

What is interesting is that all nouns which are neither sounds nor types of sounds are re-interpreted as such when selected by *listen*: media artifacts (*radio*), music artifacts (*disc*), sound makers artifacts (*bell*), events involving sound production (*cry*), speech acts (*announcement, speech*), animals (*bird*), humans (*singer, Mozart, colleague*), body parts (*chest*) and so on.

The operations at play in the various contexts, however, are different. Although they all entail re-computing (except for pure selections, as in (50a), they do not all involve the same amount of computation. For example, while (50b) involves dot exploitation, (50f) involves qualia exploitation, (50k) and (50l) involve inheritance of conventionalized attribute, and so on.

It is striking that *event* is by far the most represented type among the object arguments of *listen*. This suggests that the notion of *event* is easily re-interpretable as the type *sound*. In fact, this suggests that *sound* itself should be regarded as a type of *event*, at least in one of its possible interpretations (physical manifestation) (cf. Strawson 1990, pp. 59-86).

From a cognitive point of view, we may speculate that some shifts are easier than others: it is easier to shift from a source which is 'cognitively' close to the target than from one which is far. Conversely, source-target shifts in which the distance is bigger are cognitively more complex and less frequent.

To conclude, an exhaustive corpus analysis as proposed for *listen* shows how complex it is to classify all cases and identify the specific compositional operations at play. If we project the various contexts in (50) onto our table of prediction (cf. section 5), we can see how difficult it is to map each context into the appropriate slot. The interplay between the type system and the compositional operations seems to be more complex than the one depicted in Table 1. Notwithstanding these difficulties, we hope to have shown a that theory-informed corpus investigation as proposed here constitutes a solid methodology for a systematic description and representation of sensitivity of word meaning to context and of semantic co-compositional processes in language.

7. Concluding Observation

We have seen that the selectional behavior of words in language does indeed provide us with empirically well-founded indications of their meaning. However, the view adopted here is that a word's meaning is built from its context compositionally, and that the lexeme itself does not carry that meaning, per se. Rather, generative mechanisms in the semantics, such as coercion, modulate meanings in context and allow words to behave distributionally in unexpected ways with respect to their selectional properties. It follows that a model of the lexicon is necessary to interpret distributional data. GL seems a reasonable model for such phenomena because it provides a set of compositional rules which account for semantic processes taking place between words and phrases in text.

One of the main challenges that a theory of coercion faces, besides that of overgeneration, is the directionality of function application, since it is not always obvious what influences what in a given context. In this respect, the Head Typing Principle put forth in Asher & Pustejovsky (2006), which states that it is the syntactic head which preserves its type in composition and determines the typing of the other element(s) should be accompanied by further exploration of how multiple function application works. A related issue is how coercion phenomena and co-compositional mechanisms interact. Are they competing or collaborative principles? Finally, types prove to be insufficient to account for the whole distributional behavior of lexical items. Verbs with similar selectional properties (for instance *read* and *publish*) may exhibit different sets of collocates (cf. (8) above). Although types provide an optimal setting to capture coercion phenomena, further investigation of coercion needs to move beyond types. Further research should investigate the regularities in source-target shifts, and explore to what extent such sense modulations may occur.

Addresses of the Authors

James Pustejovsky, Dept. of Computer Science, Brandeis University Waltham, MA USA.
<jamesp@cs.brandeis.edu>

Elisabetta Jezek, Dept. of Linguistics, Pavia University, Pavia, Italy.
<jezek@unipv.it>

Notes

* We would like to thank Nicoletta Calzolari, Alessandro Lenci and Michele Prandi for comments on this paper. We would also like to thank Anna Rumshisky for her contributions to an earlier version of this work.

1 This is what is stated by the distributional hypothesis: 'difference of meaning correlates with difference in distribution' (Harris 1954, p. 156).
2 The linguistic motivations for establishing a fundamental distinction between natural and non-natural types and the conceptual underpinning of naturals are discussed in detail in Pustejovsky 2006.
3 In our analysis we use the Sketch Engine, a corpus query system which takes as input a corpus of any language (with the appropriate linguistic mark-up) and generates word sketches for the words of that language, i.e. one-page automatic, corpus-based summaries of a word's grammatical and collocational behavior (Kilgarriff et al. 2004). We use the BNC as corpus, with the following settings: minimal frequency 3, maximum number of items per grammatical relation: 150.
4 Interesting experimental work on the recognition and automatic resolution of metonymies in texts is currently under development (see for example Markert & Nissim 2006). This work, however, does not specifically address the question of how different types of metonymies can be accounted for from theoretical point of view.
5 The data below is presented adopting a layout first proposed in Rumshisky et al. (2007).
6 More exactly, the type for *house* is $phys \otimes_{telic} live\_in \bullet location$, but we will simplify for the present discussion.
7 This last case (selection of a dot object in its whole complexity) only occurs when the dot type is selected by a corresponding dot selecting predicate, like in 'read the book'.
8 It is interesting to note that Italian has another noun, *racconto* ('short story'), which has a meaning similar to that of *libro* and *romanzo* (*leggere, scrivere un racconto* 'read, write a short story'). *Racconto*, however, does not exhibit all the typical collocates of a *phys•info* type: by contrast, the presence of several verbs selecting for the *sound* dimension among its typical collocates (*ascoltare* 'listen', *sentire* 'hear', *ripetere* 'repeat', etc.) suggest that the lexical type for *racconto* is *sound•info* and that the *phys* dimension is introduced contextually by the predicator.

[9] In the literature, co-predication has been used as the main diagnostic to identify dot objects (for a definition of co-predication, see section 2 above).

[10] The combination of a verb selecting for a natural type such as *die* with an artifactual entity such as *computer* may also result in a co-composition, licensing a shift in verb meaning ((22c) below) rather than a type failure (22b) below):

(22) a. The *bird* died
b. !The *rock* died
c. My *computer* died

[11] Note that this operation does not appear in Table 1: as we clarify below, Table 1 focuses on domain-preserving coercions within the domain *entity*, while Qualia Exploitation in verb-argument contexts entails a domain-shifting coercion from *entity* to *event*.

[12] In previous GL literature, Qualia Exploitation has been discussed mostly with respect to experiencer and aspectual verbs (see for instance Pustejovsky & Bouillon 1995) but with no direct reference to the distinction between Naturals, Artifactuals and Dot types.

[13] The assumption that in constructions like (24b-c) the meaning of the verb is not affected by the differences in semantic type of the argument is not shared by Godard & Jayez (1993), who claim that in such constructions instead of type change in the argument, the semantics of the predicate is enriched to include an abstract predicate of which the complement is an argument. On the present view, the enriched interpretation is arrived at through a process of co-composition (cf. Pustejovsky 2008).

[14] It is interesting to note that naturals tend not show up in the corpus as object arguments of *finish* in its 'bring to an end' sense. This confirms the predictions of our model. Naturals are simple types with no Tensor attached: as such, they do not lend themselves to compositional operation of Qualia Exploitation, as artifactuals do. We will show, however, that naturals may participate in other kinds of compositional operations, such as Attribute Inheritance (5.1.3) and Qualia Introduction (5.2.2.).

[15] Significantly, neither natural locations like *deserts* nor natural events like *thunderstorms* tend to show up as object arguments of *attend* (see Hanks & Jezek 2008). This confirms and supports our intuition that *attend* selects for a subtype of *event*s (that is, organized events) taking place in functional locations.

[16] We assume that event introduction may be triggered not only by polymorphic predicates subcategorizing for both VP and NP complements (e.g., *finish*), but also by verbs which subcategorize exclusively for a direct object (e.g., *attend*). What is relevant is that the verb semantically selects for an event argument. On this view, we interpret syntactic subcategorization of a VP complement as syntactic evidence of the semantic selection at play.

[17] While the interpretation for 'finish the wine' in (31) results from an ordered sequence of coercions, i.e. Event Introduction and Qualia Exploitation, the interpretation for 'finish the glass' (with null complement) in (32) requires an additional operation, i.e. exploitation of the object argument of the telic value (hold (liquid)). We refer to this operation as Argument Structure Exploitation (AS-E).

[18] Several conditions have been discussed in the literature: affectedness, boundedness and modifiability of the object (Godard & Jayez 1993), aspectual restrictions on the reconstructed event (Pustejovsky & Bouillon 1995), conventionalized status of metonymic constructions (Verspoor 1997, Lapata & Lascarides 2003). For a full discussion of event coercion from a theoretical perspective, see among others Briscoe et al. (1990), Copestake & Briscoe (1995), Lascarides & Copestake (1998), Kleiber (1999), Egg (2003), Jackendoff (2002), Asher (forthcoming).

[19] Asher and Pustejovsky argue that the lexicon simplifies information that percolates up to it from commonsense metaphysics and in doing so they open up the possibility that Tensors only attach to some artifactuals and not others (Asher & Pustejovsky 2000, p. 16). In other words, by distinguishing metaphysics from the lexicon, one can both maintain that something like a door or a bathroom has a proper function without being required to claim that that function is part of the lexical entry. Although this view complicates the picture, it could partly explain why some artifacts respond easily to Qualia Exploitations and why others do not (see similar comments in Verspoor 1997, p. 189-190).
[20] Although *content* to *container* shifts like in (34) could be regarded as regular polysemy based on metonymy, we argue that coercion mechanisms such as introduction and exploitation constitute better tools for their representation than metonymic displacement. On this view, while in (34) the *container* is introduced, in (32d) the *content* is exploited.
[21] On this view, the notion of conventionalized attribute shares many similarities with the notion of "weak Quale" introduced in Busa et al. (2001).
[22] Note that in this view, (a) 'hear the alarm clock, the bell' and (b) 'hear the coffee grinder, the car' involve two different kinds of compositional operations. While in (a) the Telic Quale of the nouns is exploited, in (b) the conventional attribute of the nouns to produce noise while performing their function is inherited.
[23] For more information, see Pustejovsky (2008) for a formal analysis, and Pustejovsky & Jezek (forthcoming) for data supporting the notion of conventionalized attributes in corpus.

Bibliographical references

Asher Nicholas. To appear. A Web of Words: Lexical Meaning in Context. University of Toulouse, France and University of Texas at Austin. Ms.

Asher Nicholas & James Pustejovsky 2000. The Metaphysics of Words in Context. University of Texas at Austin and Brandeis University. Ms.

Asher Nicholas & James Pustejovsky 2006. A Type Composition Logic for Generative Lexicon. *Journal of Cognitive Science* 6. 1-38.

Bouillon Pierrette 1997. *Polymorphie et sémantique lexicale: le cas des adjectifs*. Lille: Presses Universitaires du Septentrion.

Briscoe Ted, Ann Copestake & Bran Boguraev 1990. Enjoy the Paper: Lexical Semantics via Lexicology. In Hans Karlgren (ed.). Proceedings of the 13th International Conference on Computational Linguistics (COLING-90). Helsinki. Finland. 42-47.

Busa Federica, Nicoletta Calzolari, Alessandro Lenci & James Pustejovsky 2001. Building a Semantic Lexicon: Structuring and Generating Concepts. In Blunt Harry C., Reinhard Muskens & Elias Thjisse (eds.). *Computing Meaning Vol. II*. Dordrecht: Kluwer. 29-51.

Copestake Ann, Antonio Sanfilippo, Ted Briscoe & Valeria de Paiva 1993. The ACQUILEX LKB: An Introduction. In Briscoe Ted, Valeria de Paiva & Ann Copestake (eds.). *Inheritance, Defaults, and the Lexicon*. Cambridge: Cambridge University Press.

Copestake Ann & Ted Briscoe 1995. Semi-productive Polysemy and Sense Extension. *Journal of Semantics* 12. 15-67.

CRUSE D. Alan 2000. Aspects of the micro-structure of word meanings. In RAVIN Yael & Claudia LEACOCK (eds.) *Polysemy: Theoretical and Computational Approaches*. Oxford: Oxford University Press. 30-51.

EGG Markus 2003. Beginning novels and finishing hamburgers: remarks on the semantics of *begin*. *Journal of Semantics* 20. 163-191.

FIRTH John R. 1957. A synopsis of linguistic theory 1930-1955. In *Studies in Linguistic Analysis*. 1-32. Oxford: Philological Society. Reprinted in Frank R. PALMER (ed.) *Selected Papers of J.R. Firth 1952-1959*. London: Longman (1968).

GODARD Danièle & Jacques JAYEZ 1993. Towards a proper treatment of coercion phenomena. Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics (EACL-93). Utrecht. The Netherlands. 168-177.

HANKS Patrick & Elisabetta JEZEK 2008. Shimmering Lexical Sets. In Elisenda BERNAL & Janet DE CESARIS (eds.). Proceedings of the XIII EURALEX International Congress. Barcelona. Universidad Pompeu Fabra. 391-402.

HARRIS Zellig S. 1954. Distributional structure. *Word* 10(23). 146-162.

JACKENDOFF Ray 2002. *Foundations of Language*. Oxford: Oxford University Press.

JEZEK Elisabetta & Alessandro LENCI 2007. When GL meets the corpus. A data driven investigation of semantic types and coercion phenomena. In Pierrette BOUILLON, Laurence DANLOS & Kyoko KANZAKY (eds.). Proceedings of the 4th International Workshop on Generative Approaches to the Lexicon (May 10-11, 2007). Paris. France.

KILGARRIFF Adam, Pavel RYCHLY, Pavel SMRZ & David TUGWELL 2004. The Sketch Engine. In Geoffrey WILLIAMS & Sandra VESSIER (eds.). Proceedings of the XI EURALEX International Congress (July 6-10, 2004). Lorient. France. 105-11.

KLEIBER Georges 1999. *Problèmes de sémantique. La polysémie en questions*. Lille: Presses Universitaries du Septentrion.

LAPATA Maria & Alex LASCARIDES 2003. A Probabilistic Account of Logical Metonymy. *Computational Linguistics* 29(2). 261-315.

LEVIN Beth, Grace SONG & Sue ATKINS 1997. Making sense of corpus data: a case study of verbs of sound. *International Journal of Corpus Linguistics* 2(1). 23-64.

MARKERT Katja & Malvina NISSIM 2006. Metonymic Proper Names: a Corpus-Based Account. In GRIES Stefan & Anatol STEFANOWITSCH (eds.) *Corpora in Cognitive Linguistics. Vol. 1. Metaphor and Metonymy*. Berlin & New York: Mouton-De Gruyter.

MCDONALD Scott & Michael RAMSCAR 2001. Testing the distributional hypothesis: The influence of context on judgements of semantic similarity. In Johanna D. MOORE & Keith STENNING (eds.). Proceedings of the 23rd Annual Conference of the Cognitive Science Society (August 1-4, 2001). Edinburgh. Scotland. 611-616.

MILLER George & Walter CHARLES 1991. Contextual correlates of semantic similarity. *Language and Cognitive Processes* 6(1). 1-28.

NERLICH Brigitte 2003. *Polysemy: Flexible Patterns of Meaning in Mind and Language*. Berlin & New York: Walter de Gruyter.

PRANDI Michele 2004. *The Building Blocks of Meaning: Ideas for a Philosophical Grammar*. Amsterdam & Philadelphia: John Benjamins.

PUSTEJOVSKY James 1993. Type coercion and lexical selection. In James PUSTEJOVSKY (ed.). *Semantics and the Lexicon*. Dordrecht: Kluwer Academic Publishers. 73-94.

PUSTEJOVSKY James 1995. *The Generative Lexicon*. Cambridge MA: MIT Press.

PUSTEJOVSKY James 2001. Type Construction and the Logic of Concepts. In Pierrette BOUILLON & Federica BUSA (eds.). *The Syntax of Word Meaning*. Cambridge: Cambridge University Press.

PUSTEJOVSKY James 2005. *A survey of dot objects*. Waltham Mass.: Brandeis University. Ms.

PUSTEJOVSKY James 2006. Type theory and lexical decomposition. *Journal of Cognitive Science* 6. 39-76.

PUSTEJOVSKY James 2008. *From Concepts to Meaning*. Cambridge MA: MIT Press.

PUSTEJOVSKY James & Bran BOGURAEV 1993. Lexical Knowledge Representation and Natural Language Processing. ARTIFICIAL INTELLIGENCE 63. 193-223.

PUSTEJOVSKY James & Pierrette BOUILLON 1995. Aspectual coercion and logical polysemy. *Journal of Semantics* 12(2). 133-162.

PUSTEJOVSKY James, Patrick HANKS & Anna RUMSHISKY 2004. Automated Induction of Sense in Context. Proceedings of the 20th International Conference on Computational Linguistics (COLING-04). Geneva. Switzerland. 924-931.

PUSTEJOVSKY James & Elisabetta JEZEK Forthcoming. *Meaning in Context. An Introduction to Generative Lexicon Theory*. Ms.

PUSTEJOVSKY James & Anna RUMSHISKY 2008. Between Chaos and Structure: Interpreting Lexical Data Through a Theoretical Lens. *International Journal of Lexicography* 21. 337-355.

RUMSHISKY Anna, Victor A. GRINBERG & James PUSTEJOVSKY 2007. Detecting selectional behavior of complex types in text. Proceedings of the 4th International Workshop on Generative Approaches to the Lexicon (May 10-11, 2007). Paris. France.

RUMSHISKY Anna 2008. Resolving polysemy in verbs: Contextualized distributional approach to argument semantics (this volume).

SAHLGREN Magnus 2006. *The Word-Space Model. Using distributional analysis to represent syntagmatic and paradigmatic relations between words in high-dimensional vector spaces*. Stockholm University. PhD dissertation.

STRAWSON Peter Frederick. 1990. *Individuals: An essay in descriptive metaphysics*. London & New York: Routledge [First ed. 1959].

VERSPOOR Cornelia M. 1997. *Contextually-dependent Lexical Semantics*. University of Edinburgh. PhD dissertation.

Chapter 3

Co-Compositionality. *James Pustejovsky*

Chapter 17
Co-compositionality in Grammar

James Pustejovsky

Abstract

This entry addresses the problem of how words combine to make meanings that appear non-compositional in their derivation. Specifically, we examine a phenomenon known as co-compositionality, where a new meaning emerges in an expression that is not expected through simple compositional operations. This will be analyzed as a kind of bilateral function application, where both predicate and argument contribute functionally to determine the meaning of the resulting expression. In this article, we differentiate the formal properties of co-compositionality from conventional mechanisms of composition, and examine two co-compositional constructions at work in language: (a) cospecification, where the argument to a verb acts functionally over the predicate selecting it; and (b) subject-induced coercion, where the subject adds an agentive or intentional interpretation to the meaning of the predicate selecting it. All cases of co-composition are ampliative, in that the meaning of the derived expression entails the meanings of the subexpressions. By studying the mechanisms of such constructions, we hope to arrive at a better understanding of the mechanics of argument selection, and with this, a richer appreciation for the nature of compositionality in language.

1 Basic Mechanisms of Selection

Co-compositionality is a semantic property of a linguistic expression in which all constituents contribute functionally to the meaning of the entire expression. As a result, it extends the conventional definition of compositionality. The principle of compositionality in linguistics (cf. Janssen, 1983, Thomason, 1974) and in philosophy (cf. Werning, 2004) involves the notion that the meanings of complex symbols are systematically determined by the composition of their component parts. In order to understand the theoretical motivation behind the theory of co-compositionality, it is necessary to understand where conventional theories of compositionality are unable to explain the meaning of certain natural language constructions. Since these issues are addressed in more detail by other entries in this encyclopedia (Non-compositionality), the present article will focus on the role that compositionality plays in mapping from the lexicon to syntactic form.

At the outset, it should be stated that co-compositonality is not the result of a failure of compositionality, and hence to be viewed as involving non-compositional processes. Rather, as the name would suggest, it entails at least conventional compositional mechanisms for the expressions involved, along with additional interpretive mechanisms not always exploited within a phrasal composition. In order to understand what these are, we first review conventional modes of argument selection in language.

While it is impossible to say how many meanings we create for a particular word in normal language use, we can reasonably ask how many meanings we have stored for that word in our mental lexicon. This is where linguists differ broadly in assigning responsibility for whether meaning shifts occur at all and, if so, how. As a result of this divide, the role that compositionality plays in structuring not only the grammar but also the lexicon is significant.

For example, in conventional models of language meaning, a verb is thought to have several different word senses. For each sense, the verb acts on its parameters (its arguments in syntax) in a compositional manner. This means that the semantics of the result of application of the verbal function to its argument is determined by the semantics of the function itself, a process referred to as *function application*. Consider, for example, the way in which the verbs *throw* and *kill* each have several distinct senses.

(1) a. Mary threw the ball to John. (PROPEL)
b. They threw a party for Bill. (ORGANIZE)
c. Mary threw breakfast together quickly. (CREATE)

The use of *throw* in each sentence above illustrates a true verbal ambiguity, one that requires separate senses, each with specific subcategorization and semantic selection as illustrated. Likewise, the verb *kill* as used in (2) below, demonstrates a systematic sense distinction as well.

(2) a. John killed the plant.
b. Mary killed the conversation.
c. John killed the evening watching TV.

As with the verb *throw*, each of these senses has a regular and productive distribution in the language, exemplified below.

(3) a. Mary killed the fish.
b. The President killed any attempt at dialogue with Cuba.
c. John killed the day reading.

Verb senses like these are distinct, semantic units, perhaps related to each other, but stored separately in the lexicon. Because they have distinct subcategorization and type selection frames, the semantic computation involving these senses in the syntax can be performed compositionally.

These examples with the verbs *throw* and *kill* illustrate that lexical forms may be truly ambiguous, and as such, can be modeled adequately by a sense enumerative lexical (SEL) model (cf. Pustejovsky, 1995). In such a model, each sense of a word, as in (2) above, would be strongly typed, illustrated in (4) below, where the intended sense is glossed as a relation with its appropriate argument types.

(4) a. kill_1: CAUSE-TO-DIE(THING, ANIMATE)
b. kill_2: TERMINATE(HUMAN, EVENT)
c. kill_3: SPEND(HUMAN, TIME, EVENT)

Given distinct lexical types for these three senses of *kill*, compositional mechanisms in the semantics can compute the sentences in (2) as cases of function application. For this particular example, *function application* assumes that the verb *kill* applies to its arguments in discrete steps. For example, consider the derivation of (2c) as a sequence of function applications, simplifying the arguments (HUMAN, TIME, EVENT) from (4c) as numbered variables.

(5) a. John killed the day reading.
b. $kill(Arg_1, Arg_2, Arg_3)$
c. Apply $kill(Arg_1, Arg_2, Arg_3)$ to "reading"
$\Longrightarrow kill(Arg_1, Arg_2, [reading])$
d. Apply $kill(Arg_1, Arg_2, [reading])$ to "the day"
$\Longrightarrow kill(Arg_1, [day], [reading])$
e. Apply $kill(Arg_1, [day], [reading])$ to "John"
$\Longrightarrow kill(john, [day], [reading])$

This derivation has a successful computation because the verb sense for *kill* selected in (5) has the appropriate typing. If we had tried using the type associated with kill_2, the sentence would not have an interpretation. As we see, compositional operations reflect the ontological and lexical design decisions made in the grammar.

Treating the functional behavior of composition formally, we can state this procedure as an operation over the types of expressions involved, as expressed in (6):

(6) Function Application (FA):
If α is of type a, and β is of type $a \to b$, then $\beta(\alpha)$ is of type b.

Returning to the example derivation in (5), we can see FA at work on the last application step in (5e), where e stands for any of the specific types mentioned earlier (e.g., THING, HUMAN, TIME, EVENT) and t stands for the propositional type.

(7) a. $kill(Arg_1, [day], [reading])$ is of type $e \to t$;
b. $john$ is of type e;
c. FA results in applying $e \to t$ to e;
$\Longrightarrow kill([john], [day], [reading])$, of type t, i.e., a sentence.

Hence, by enumerating separate senses for ambiguous predicates, we can ensure strong (unique) typing on the arguments expected by a verb (function), and thereby maintain compositionality within these constructions.

If function application as described above were inviolable, then we would not expect to encounter examples of type mismatch between verb and argument. But, of course, such data are ubiquitous in language, and involve a process characterized as *type coercion* (Pustejovsky, 1995, Copestake and Briscoe, 1995, Partee and Rooth, 1985). This is an operation that allows an argument to change its type, if it does not match the type requested by the verb. For example, for one of its senses, the aspectual verb *begin* selects for an event as its internal argument:

(8) Mary began [reading the book]$_{event}$.

The same sense is used, however, when *begin* selects for a simple NP direct object, as in (9).

(9) Mary began [the book]$_{event}$.

In such configurations, the verb is said to "coerce" the NP argument into an event interpretation (cf. Pustejovsky, 1991,1995). Under such an analysis, the NP actually denotes a salient event that involves the book in some way, e.g., reading it, writing it, and so on. This is schematically represented below, where the NP *the book* has been reinterpreted through coercion, as some relation, *R*, involving the book.

(10)

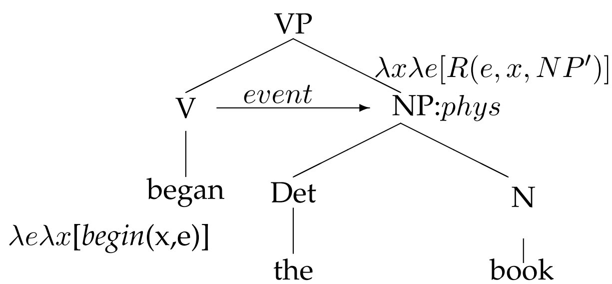

Our knowledge of the world associates conventional activities, such as reading and writing, with books. This knowledge can be lexically encoded through the use of *Qualia Structure* (Pustejovsky, 1995), thereby providing a mechanism for preserving compositionality in the construction above. In Generative Lexicon Theory (Pustejovsky, 1995), it is assumed that word meaning is structured on the basis of four generative factors (the Qualia roles) that capture how humans understand objects and relations in the world and provide the minimal explanation for the linguistic behavior of lexical items (these are inspired in large part by Moravcsik's (1975, 1990) interpretation of Aristotelian *aitia*). These are: the FORMAL role: the basic category that distinguishes the object within a larger domain; CONSTITUTIVE role: the relation between an object and its constituent parts; the TELIC role : its purpose and function; and the AGENTIVE role: factors involved in the object's origin or "coming into being". Qualia structure is at the core of the generative properties of the lexicon, since it provides a general strategy for creating new types.

The qualia act as type shifting operators, that can allow an expression to satisfy new typing environments. Every expression, α, has some set of operators available to it, that provide such type shifting behavior. Let us refer to this set as Σ_α. Then we can characterize function application under such conditions as follows:

(11) Function Application with Coercion (FA$_c$):
If α is of type c, and β is of type $a \rightarrow b$, then,
(i) if type c = a then $\beta(\alpha)$ is of type b.
(ii) if there is a $\sigma \in \Sigma_\alpha$ such that $\sigma(\alpha)$ results in an expression of type a, then $\beta(\sigma(\alpha))$ is of type b.
(iii) otherwise a type error is produced.

Such phenomena are quite common in language, and when viewed as a lexically-triggered operation, coercion allows us to maintain a compositional treatment of argument selection in the grammar.

2 Co-compositional Mechanisms

With the additional mechanism of function application with coercion (FA_c), we are able to account for a larger range of data that would otherwise not have been modeled as compositional in nature. But there are many constructions in language which appear to be outside the scope of conventional compositional operations. In this section, we see how these can be analyzed co-compositionally.

As stated above, co-compositionality is a semantic property of a linguistic expression in which all constituents contribute functionally to the meaning of the entire expression. As with compositionality, the notion of co-compositionality is a characterization of how a system constructs the meaning from component parts. It is a mistake to think that an expression in a language is inherently co-compositional or compositional. Rather, it is the set of computations within a specific system that should be characterized as co-compositional for those expressions. To make this distinction clear, consider the verb *run* as it is used in the contexts of (12)-(13) below.

(12) a. John ran.
b. John ran for twenty minutes.
c. John ran two miles.

(13) a. John ran to the store.
b. John ran the race.

There are two senses of *run* that emerge in context with these examples:

(14) a. run_1: manner-of-motion activity, as used in (12);
b. run_2: change-of-location transition, as used in (13);

We can choose to design our semantics and the accompanying lexicon for these cases according to the null hypothesis, and create separate senses, as illustrated in (14). With two separate entries, they will select differently because they will have different types and argument structures. In this case, we say that the data are accounted for compositionally through sense enumeration. What is left unexplained, however, is any logical relation between the senses, a major drawback; this can be overcome, however, with lexical rules that explicitly specify this relationship as a redundancy rule or meaning postulate.

Similar remarks hold for verbs such as *wax* and *wipe* in (15)–(16), which are contextually ambiguous between a process reading and a transition reading, depending on the presence of a resultative adjectival. Normally, lexicons would have to enter both forms as separate lexical entries (cf. Levin and Rappaport, 1995).

(15) a. Mary *waxed* the car.
b. Mary *waxed* the car clean.

(16) a. John *wiped* the counter.
b. John *wiped* the counter dry.

Clearly, the local context is supplying additional information to the meaning of the predicate that is not inherently part of the verb's meaning; namely, the completive aspect that inheres in the resultative constructions (cf. Goldberg (1995) and Jackendoff (2002)).

A related phenomenon of extended word sense in context is what Atkins et al (1988) refer to as "overlapping senses", and it is exhibited by cooking verbs such as *bake*, *fry*, as well as by activities such as *carve*, shown below.

(17) a. John *baked* the potato.
b. John *baked* the cake.

(18) a. Mary *fried* an egg.
b. Mary *fried* an omelette.

(19) a. John *carved* the stick.
b. John *carved* a statue.

These example illustrate that strict lexical typing (preserving compositionality) does not explain when and how verb senses will overlap or be entailed by another sense. Clearly, something is not being captured by the semantic theory with such data. The notion of co-compositionality was introduced to characterize just this type of phenomenon (cf. Pustejovsky, 1991, 1995), In particular, this construction has been referred to as *cospecification*, since the argument being selected by the predicate, seems to have a semantic familiarity with the predicate, and hence, *specifies* the governing predicate.

Informally, we can view co-compositionality as the introduction of new information to an expression by the argument, beyond what it contributes as an argument to the function within the phrase. Hence, it can be considered an *ampliative* operation, relative to the function application. Returning to the examples considered above, let us see how this characterization fits the data. First, consider the shift from the process interpretation of *run* to the accomplishment sense in (12)-(13). The sense of the verb *run* in (13b) clearly overlaps (indeed, it entails) the sense exploited in (12a). We say that the NP *the race* in (13b) cospecifies the predicate selecting it, repeated below in (20).

(20) John ran the race.

The semantic composition results in an interpretation entailing the activity of running, which is either quantified by a measure phrase with a specific distance (as in (12c) with *two miles*), or entails the completion of a specific course or event (as in (20) with *the race*).

With the verbs *wax* and *wipe*, similar extensions to the basic meaning are at play in (15b) and (16b). What is still unclear is how the extended meaning is first licensed and then how it is computed formally through compositional mechanisms.

To better understand the mechanisms involved in the ampliative interpretations that result in such constructions, we examine the relationship between the core and derived senses of the verb *bake*, as presented above in (17). In the context of particular objects, the verb *bake* assumes the interpretation of a *creation* predicate, while with other objects, it maintains the underlying *change-of-state* predicate meaning. Certain NPs are said to *cospecify* the verb selecting it, as does the noun *cake* in its agentive qualia value. That is, the type structure for *cake* references the predicate selecting it as an argument. With this, the activity of baking assumes a resultative interpretation when combined with co-specifying arguments.

Assume that the lexical semantics for the change-of-state sense of *bake* is given as in (21), where the qualia roles are abbreviated as F (Formal), C (Constitutive), T (Telic), and A (Agentive).

(21) $$\lambda y \lambda x \lambda e \begin{bmatrix} \textbf{bake} \\ \text{AS} = \begin{bmatrix} \text{A}1 = x : phys \\ \text{A}2 = y : phys \end{bmatrix} \text{ES} = \begin{bmatrix} \text{E}1 = e : process \end{bmatrix} \\ \text{QS} = \begin{bmatrix} \text{A} = bake\_act(e, x, y) \end{bmatrix} \\ \end{bmatrix}$$

The lexical representation for an artifactual concept such as the noun *cake* is shown below in (22).

(22) $\lambda x \exists y \left[\begin{array}{l}\textbf{cake} \\ \text{AS} = \left[\begin{array}{l}\text{ARG1} = x : phys \\ \text{D-ARG1} = y : mass\end{array}\right] \text{QS} = \left[\begin{array}{l}\text{F} = cake(x) \\ \text{C} = made\_of(x,y) \\ \text{T} = \lambda z, e[eat(e,z,x)] \\ \text{A} = \exists w, e[bake(e,w,y)]\end{array}\right]\end{array}\right]$

Notice that the Agentive qualia value for the noun *cake* makes reference to the very process within which it is embedded in the sentence in (17) (i.e., *bake a cake*), which is a case of cospecification.

We now define the conditions under which the derivation of an expression is said to be *co-compositional*. Ignoring the event structure for discussion, according to the type structure for the predicate *bake*, function application, as defined above, applies as expected to its argument *a cake*.[1] But the direct object cospecifies the verb selecting it, since its type structure makes reference to the governing verb, *bake*. This is illustrated graphically in (23).

(23)

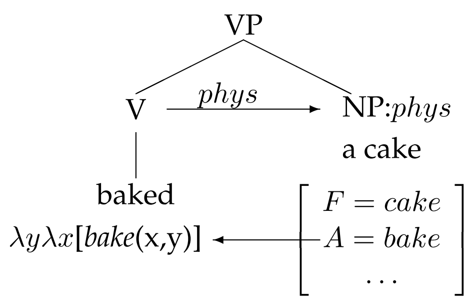

From the underlying process *change-of-state* sense of *bake*, the *creation* sense emerges when combined with the triggering NP *a cake*. This results in a logical form such as that shown in (24).

(24) $\exists e_1 \exists e_2 \exists x \exists y[bake(e_1, j, y) \wedge cake(e_2, x) \wedge made\_of(x, y) \wedge e_1 \leq e_2]$

The operation of co-composition results in a qualia structure for the VP that reflects aspects of both constituents. These include:

(A) The governing verb *bake* applies to its complement;

(B) The complement co-specifies the verb;

[1] We also ignore the type shifting involved for the predicate to take the generalized quantifier *a cake* as its argument. For discussion, we assume the indefinite is treated as a discourse variable denoting an individual type.

(C) The composition of qualia structures results in a derived sense of the verb, where the verbal and complement AGENTIVE roles match, and the complement FORMAL quale becomes the FORMAL role for the entire VP.

The derived sense is computed from an operation called *qualia unification*, introduced in Pustejovsky (1995). The conditions under which this operation can apply are stated in (25) below:

(25) FUNCTION APPLICATION WITH QUALIA UNIFICATION: For two expressions, α, of type `<a,b>`, and β, of type `a`, with qualia structures QS_α and QS_β, respectively, then, if there is a quale value shared by α and β, $[_{QS_\alpha} \ldots [Q_i = \gamma\ ]]$ and $[_{QS_\beta} \ldots [Q_i = \gamma\ ]]$, then we can define the qualia unification of QS_α and QS_β, $QS_\alpha \sqcap QS_\beta$, as the unique greatest lower bound of these two qualia structures. Further, $\alpha(\beta)$ is of type `b` with $QS_{\alpha(\beta)} = QS_\alpha \sqcap QS_\beta$.

The composition in (23) can be illustrated schematically in (26) below.

(26) $\begin{bmatrix}\text{V} & A = bake\end{bmatrix} \sqcap \begin{bmatrix}\text{NP} & \begin{matrix} F = cake \\ A = bake \end{matrix}\end{bmatrix} = \begin{bmatrix}\text{VP} & \begin{matrix} F = cake \\ A = bake \end{matrix}\end{bmatrix}$

3 Further Extensions of Co-composition

Further examination of the derivation above suggests that co-composition involves a more general process where conventional function application from an anchor function (e.g., the governing verb), along with ampliative information supplied by a triggering argument type. These properties can be summarized as follows in (27).

(27) Properties of Co-compositional Derivations:
a. Within an expression, α, consisting of two subexpressions, α_1 and α_2, i.e., $[_\alpha\ \alpha_1\ \alpha_2]$, one of the subexpressions is an *anchor* that acts as the primary functor;
b. Within the argument expression, there is explicit reference to the anchor or the anchor's type (that is, the complement co-specifies the functor);
c. The composition of lexical structures results in a derived sense of the functor, within α.

This can be formalized as follows:

(28) Co-compositionality:
a. The derivation for an expression α, is *co-compositional* with respective to its constituent elements, α_1 and α_2, if and only if one of α_1 or α_2 applies to the other, $\alpha_i(\alpha_j)$, $i \neq j$, and $\beta_j(\alpha_i)$, for some type structure β_j within the type of α_j, i.e., $\beta_j \sqsubseteq type(\alpha_j)$.
b. $[\![\alpha]\!] = \alpha_i(\alpha_j) \sqcap \beta_j(\alpha_i)$.

For the example at hand, the overall expression α is *bake a cake*. The anchor functional term is the verb *bake* (α_1), and the ampliative interpretation comes from the Agentive Qualia value for the NP (β_j). Given this formulation of co-composition, it is now clear now when co-composition is licensed. If any component of the type of the argument in a construction makes reference to the anchor functional term in a construction, then co-composition should be permitted. This is, in fact, what we see in all the cases of cospecification we encountered above.

With the more general characterization of composition given above, we can now analyze a number of constructions as co-compositional in nature. These include, among others, subject-derived agentive interpretations (*subject-induced coercion*) and certain light verb constructions, e.g., *functionally dependent verbs*. For example, it has long been noted that certain classes of predicates select for non-agentive subjects, but allow agentive interpretations in the appropriate context, as illustrated in the examples below (cf. Wechsler, ref, others).

(29) a. The storm killed the deer.
b. An angry rioter killed a policeman.

(30) a. The glass touched the painting.
b. The curious child touched the painting.

(31) a. The ball rolled down the hill.
b. John rolled down the hill as fast as he could.

(32) a. The room cooled off quickly.
b. John cooled off with an iced latte.

We will refer to these as *subject-induced coercions*, since, in each of these pairs, the subject in the (b)-sentence introduces agency or intentionality towards the predicated event. Rather than suggesting that each of these verbs is ambiguous between agentive/non-agentive readings, we can view the computation in the (b)-sentences as co-compositional, where an agentive subject introduces the appropriate intentional component to the interpretation of the VP. For the present discussion, let us characterized "agency", in terms of Qualia Structure, as referring to the potential to act towards a goal. For a cognitive agent, such as a *human*, this amounts to associating a set of particular activities, $\mathcal{A}$, as the value of the Agentive role, and a set of goals, $\mathcal{G}$, associated with the Telic role in the Qualia for that concept, as illustrated below in (33).

(33) $$\lambda x \begin{bmatrix} \textbf{human\_agent} & \\ \text{QS} = & \begin{bmatrix} \text{F} = human(x) \\ \text{T} = \lambda e'[\mathcal{G}(\mathbf{e',x})] \\ \text{A} = \lambda e[\mathcal{A}(\mathbf{e,x})] \end{bmatrix} \end{bmatrix}$$

Consider how this composition is instantiated for the subject-induced coercion in (29b). Causative verbs such as *kill* denote transitions from one state to a resulting state, by virtue of a causing event. This can be represented as the lexical structure given in (34).

(34) $$\lambda y \lambda x \lambda e_2 \lambda e_1 \begin{bmatrix} \textbf{kill} & & & \\ \text{AS} = & \begin{bmatrix} \text{A1} = x : phys \\ \text{A2} = y : phys \end{bmatrix} & \text{ES} = & \begin{bmatrix} \text{E1} = e_1 : process \\ \text{E2} = e_2 : state \end{bmatrix} \\ \text{QS} = & \begin{bmatrix} \text{F} = dead(e_2, y) \\ \text{A} = kill\_act(e, x, y) \end{bmatrix} & & \end{bmatrix}$$

Co-composition of the subject with the VP results in an agentive predicate replacing the underspecified predicate (i.e., $kill\_act$) in the VP's agentive Qualia Structure. The resulting interpretation is shown in (35).

(35) $$\exists x, y, e_1, e_2[rioter(e_1, x) \wedge \mathcal{A}(e_1, x, y) \wedge police(y) \wedge dead(e_2, y) \wedge e_1 \leq e_2]$$

In fact, most cases of subject-induced coercion can be characterized in the manner defined above, as ampliative readings resulting from co-composition (cf. Pustejovsky, 2011 for further discussion).

Another interesting case of co-composition can be seen in certain light verb constructions (Rosen, 1997, Goldberg, 1995, Butt, 1997, Mohanan, 1997), where much of the semantic content of the predicate is contributed by the complement meaning. Of particular interest to the current discussion are *functionally dependent verb* readings (Pustejovsky, 1995). These involve a range of verb classes, characterized by the verb's dependence on the specific *function* of the complement selected. Included in this class are the verbs *open*, *close*, *break*, and *fix*. The problem for compositionality for light verb constructions in general, and this class in particular, is the recurring issue of sense specificity. That is, can the different uses of *open*, for example, in (36) , be captured with one verb meaning or are multiple senses required?

(36) a. Mary opened the letter from her mother.
b. The rangers opened the trail for the season.
c. John opened the door for the guests.
d. Mary opened up the application.
e. She then opened a window and started writing.

Viewed as a co-compositional operation, in each case above, the sense of the verb *open* has been enriched through the context of the meaning associated with a specific object type. As with subject-induced coercions, the resulting VP meaning is ampliative relative to the function application of the verb over its object. This additional inference is derived from the complement itself. Briefly, we can view the verb *open* as bringing about a change of state, one which *enables* the activities associated with the complement's TELIC role. These are spelled out, somewhat informally in the glosses for each of the cases in (36) below.

(37) a. The letter can now be *read*.
b. The trail can now be *walked on*.
c. The door can be *walked through*.
d. The application is *running*.
e. The window is *ready for typing*.

1 Future Directions

In this entry we have defined the general characteristics associated with co-compositional analyses of a modest range of linguistic phenomena. It is obvious that there is much still to study with the behavior of co-compositionality in language. For example, there are clearly *degrees* of co-compositionality in the cases we have reviewed, and even more with cases we have not presented here. Current research on these areas focus on broadening the definition of co-composition to include both finer degrees of sense modulation (cf. Pustejovsky and Rumshisky, 2009, Pustejovsky and Jezek, 2008), and deeper sense extensions to metaphorical shifts of meaning (cf. Pustejovsky and Rumshisky, 2010).

References

Apresjan, J. D. 1973a. “Synonymy and Synonyms,” in F. Kiefer (ed.), 173-199.
Borschev, Vladimir and Barbara H. Partee. Genitive modifiers, sorts, and metonymy. Nordic Journal of Linguistics, 24(2):140?160, 2001.
Bouillon, P. and F. Busa (eds) 2001. *The Language of Word Meaning*, Cambridge University Press, Cambridge.
Butt, M. 1997. Complex Predicates in Urdu. In A. Alsina, J. Bresnan, and P. Sells eds. Complex Predicates. Stanford, California: CSLI Publications, p. 107-149.
Comrie, Bernard. 1981.Language universals and linguistic typology: Syntax and morphology. Chicago: University of Chicago Press.
Copestake, A. 1992.*The Representation of Lexical Semantic Information*, CSRP 280, University of Sussex.
Copestake, A. (1993) “Defaults in the LKB”, in T. Briscoe and A. Copestake (Eds.) *Default Inheritance in the Lexicon*, Cambridge University Press.
Copestake, A. and E. Briscoe. (1992) “Lexical Operations in a Unification-Based Framework,” in J. Pustejovsky and S. Bergler (Eds.) *Lexical Semantics and Knowledge Representation*, Springer Verlag, New York.
Copestake, A. and T. Briscoe, 1995. Semi-productive Polysemy and Sense Extension. In *Journal of Semantics*, 15-67.
Davidson, D. 1967. “The Logical Form of Action Sentences,” in N. Rescher, ed., The Logic of Decision and Action, Pittsburgh University Press, Pittsburgh.
Dölling, J. (1992) “Flexible Interpretationen durch Sortenverschiebung”, in Ilse Zimmermann, Anatoli Strigen (eds.) *Fügungspotenzen*, Berlin, Akademie Verlag.
Dowty, D. R. 1979. *Word Meaning and Montague Grammar*, D. Reidel, Dordrecht, Holland.
Dowty, D. R. (1986). The effects of aspectual class on the the temporal structure of discourse: Semantics or pragmatics, Linguistics and Philosophy, 9.1
Dowty, D. (1991). Thematic proto-roles and argument selection. Language 67:547-619.
Goldberg, A. E. (1995). Constructions: A Construction Grammar Approach to Argument Structure. Chicago: University of Chicago Press.
Grice, H. P. (1969). ”Utterer’s Meaning and Intentions”. Philosophical Review, 78 (1969), 147-177. In Grice (1989).
Grimshaw, J., and A. Mester. 1988. Light verbs and theta-marking. Linguistic Inquiry 19(2):205232.
Jackendoff, R. (1972). Semantic Interpretation in Generative Grammar. Cambridge, MA: MIT Press.
Jackendoff, R. 1990. *Semantic Structures*, MIT Press, Cambridge, MA.
Jackendoff, R.S. 1992. “Babe Ruth Homered His Way into the Hearts of America,” in T. Stowell and E. Wehrli (eds.), Syntax and the lexicon, Academic Press, San Diego, 155-178
Jackendoff, R. 1997. *The Architecture of the Language Faculty*, MIT Press, Cambridge, MA.
Jackendoff, R. 2002. *Foundations of Language*, Oxford University Press, Oxford.
Janssen, T. M. V. (1983) Foundations and Applications of Montague Grammar. Amsterdam: Mathematisch Centrum.
Janssen, T. M. V. (1997) “Compositionality” In J. van Benthem and A. ter Meulen eds., Handbook of Logic and Language. Amsterdam: Elsevier. 417?473.
Krifka, M. 1992. “Thematic Relations as Links between Nominal Reference and Temporal Constitution,” in I. Sag and A. Szabolcsi, (eds.), *Lexical Matters*, CSLI Lecture Notes, University of Chicago Press, Chicago, Ill.

Levin, B., and M. Rappaport Hovav (1995). Unaccusativity: at the Syntax-Semantics Interface. Cambridge, MA: MIT Press.
Levin, B., and M. Rappaport Hovav (2005). Argument Realization. Cambridge University Press. Moens, M. and M. Steedman. 1988."Temporal Ontology and Temporal Reference," *Computational Linguistics* 14:15-28.
Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In Jaakko Hintikka, Julius Matthew Emil Moravcisk, and Patrick Suppes, eds., Approaches to Natural Language, 221?242. Dordrecht: D. Reidel. Reprinted in Montague (1974), 247?270.
Montague, Richard. 1974. Formal Philosophy: Selected Papers of Richard Montague. New Haven, CT: Yale University Press.
Moravcsik, J. M. 1975. "Aitia as Generative Factor in Aristotle's Philosophy", *Dialogue*, 14:622-36.
Parsons, T. 1990. Events in the Semantics of English, MIT Press, Cambridge, MA.
Partee, Barbara H. 1984. Compositionality. In Fred Landman and Frank Veltman, eds., Varieties of Formal Semantics, 281?311. Dordrecht: Foris.
Partee, B. and M. Rooth. (1983) "Generalized Conjunction and Type Ambiguity", in *Meaning, Use, and Interpretation of Language*, Bäuerle, Schwarze, and von Stechow (eds). Walter de Gruyter.
Pustejovsky, J. 1991a. "The Generative Lexicon," *Computational Linguistics*, 17:409-441.
Pustejovsky, J. 1991b. "The Syntax of Event Structure," *Cognition* 41:47-81.
Pustejovsky, J. (1995) *The Generative Lexicon*, MIT Press, Cambridge, MA.
Pustejovsky, J. (2011). *The Multiplicity of Meaning*, MIT Press, Cambridge.
Pustejovsky, J. and E. Jezek (2008). "Semantic Coercion in Language: Beyond Distributional Analysis," *Italian Journal of Linguistics* 20:1, 181-214.
Pustejovsky, J. and A. Rumshisky. 2008. Between Chaos and Structure: Interpreting Lexical Data through a Theoretical Lens, In International Journal of Lexicography, OUP.
Pustejovsky, J. and A. Rumshisky (2010). "Mechanisms of Sense Extensions in Verbs", in Gilles-Maurice de Schryver (ed.) *A Way with Words: A Festscrhift for Patrick Hanks*, Menha Publishers.
Tenny, C. 1992. "The Aspectual Interface Hypothesis," in I. A. Sag and A. Szabolcsi (eds.), 1-27.
Tenny, C. and J. Pustejovsky (2000) *Events as Grammatical Objects*, CSLI Publications, Stanford, University of Chicago Press.
Rosen, C. 1997. Auxiliation and Serialization: On Discerning the Difference. In A. Alsina, J. Bresnan, and P. Sells eds. Complex Predicates. Stanford, California: CSLI Publications, p. 175- 202.
Thomason, R. (1974) (ed.) Formal Philosophy. Selected Papers by Richard Montague. New Haven.
Weinreich, U. *Explorations in Semantic Theory*, Mouton, The Hague, 1972.
Werning, M. (2004). "Compositionality, Context, Categories and the Indeterminacy of Translation" *Erkenntnis* 60 (2).

Chapter 4

Verb Meaning in Context: Integrating VerbNet and GL Predicative Structures. *James Pustejovsky, Martha Palmer, Annie Zaenen, and Susan Brown*

Verb Meaning in Context:
Integrating VerbNet and GL Predicative Structures

James Pustejovsky[†], Martha Palmer[‡], Annie Zaenen[§], and Susan Brown[‡]
[†]Brandeis University [‡]University of Colorado [§]Stanford University
jamesp@cs.brandeis.edu, {martha.palmer,susan.brown}@colorado.edu, azaenen@stanford.edu

Abstract

This paper reports on aspects of a new research project aimed at enriching VerbNet's predicative structures with representations and mechanisms from Generative Lexicon Theory. This involves the introduction of systematic predicative enrichment to the verb's predicate structure, including an explicit identification of the mode of opposition structure inherent in the predicate. In addition, we explore a GL-inspired semantic componential analysis over VerbNet classes, in order to identify coherent semantic cohorts within the classes.

Keywords: Event Semantics, Event Structure, VerbNet, Generative Lexicon

1. Introduction

In this research note, we report on a newly funded effort towards integrating VerbNet's lexical structure and Generative Lexicon's (GL) semantic representation.[1] Our overall goal is to address two of the major problems in the representation and annotation of verb meaning in natural language: (i) how to encode the context-dependence of the meaning of a verb; and (ii) how to adequately represent the subeventual predication that inheres in complex verb meanings and is associated with polysemy arising in distinct contexts. Specifically, we propose integrating GL's compositional approach to event semantics with the predicative representations in VerbNet. This includes making explicit reference to the conditions that hold before, during, and as a result of an activity or event. Here we focus primarily on the predicative content of these conditions and how this technique might contribute additional structural distinctions within VerbNet classes.

It is well known that verbs can be notoriously polysemous. Sometimes this occurs with overt syntactic markers that are relatively easy to identify, as when a "moved" argument alternation signals both a new subcategorization frame as well as a shift in meaning, as illustrated in (1) below. In fact, there is controversy over whether such meaning preserving diathesis alternations actually constitute true polysemy or not (Levin, 1993).[2]

(1) a. *The wind broke the glass.*
break-45.1, [NP V NP]
b. *The glass broke suddenly.*
break-45.1, [NP.patient V]

But just as often, polysemy emerges not from argument alternation, but from PP or other forms of predicative adjunction, cf. (2).

(2) a. *The books slid.*
slide-11.2, [NP V]
b. *The books slid from the table.*
slide-11.2, [NP V PP.init_loc]
c. *The books slid to the floor.*
slide-11.2, [NP V PP.dest]

Here we see a manner-of-motion verb lexically typed as a process in (2a), and in (2b)and (2c) as a telic event. The semantics for each of these senses is illustrated in (3):[3]

(3) a. **[NP V]**: motion(during(E), Theme)
b. **[NP V PP.init_loc]**: motion(during(E), Theme) path_rel(start(E), Theme, Init_Loc, ch_of_loc, prep)
c. **[NP V PP.dest]**: motion(during(E), Theme) path_rel(end(E), Theme, Dest, ch_of_loc, prep)

Other examples can be seen with the verbs *yank* and *push*.

(4) a. *Nora yanked the button loose.*
push-12-1, [NP V NP ADJP-Result]
b. *Nora pushed the tables apart.*
push-12-1, [NP V NP ADJP-Result]

These are typically analyzed as cases of constructional meaning (Goldberg and Jackendoff, 2004; Croft, 2001) or co-composition (Pustejovsky, 1995b; Pustejovsky and Busa, 1995), where the construction associated with these examples reflects a contextualized interpretation of the verb meaning. For example, the semantic representation for *yank* in (4a) given in VerbNet is shown in (5).

(5) cause(Agent, E) contact(during(E), Agent, Theme) exert_force(during(E), Agent, Theme) Pred(result(E), Theme)

Further, notice that the two verbs in (4) are currently annotated as members of the same VerbNet class, since the goal

[1]This work is being carried out in the context of two grants: CwC, a DARPA effort to identify and construct computational semantic elements, for the purpose of carrying out joint plans between a human and computer through NL discourse: and eTASC, a DTRA effort to identify and build semantic components in natural language.

[2]In the examples below, we annotate verb uses with VerbNet class identifiers and the specific construction invoked (Kipper et al., 2006; Brown et al., 2014).

[3]See (Hwang et al., 2014) for discussion of path_rel in the context of motion and caused motion constructions.

of VerbNet is to capture commonalities of syntax-semantics interaction across members of a class. However, this leaves within-class semantic distinctions still needing further clarification.

The representations make no mention of predicative content differentiating them from other within-class verb members (such as *bounce* and *float*). In addition, the approach VerbNet currently uses for capturing event structure, which distinguishes between the start, end and middle (during) of an event, does not always provide a consistent, detailed representation of different event structures for different types of events.

From such observations, we have begun exploring how the Qualia and Event Structures from Generative Lexicon Theory (GL) can help overcome some of these problems. First, by incorporating a richer subeventual predicative structure within VerbNet's representation, we will be better able to distinguish within-class coherent semantic groupings. Secondly, a more structured and compositional approach to subeventual semantics will help explain the semantics encoded in cases of VerbNet constructional polysemy. In the remainder of this short note, we focus on how to enrich VerbNet's predicative structure, while deferring discussion of changes to the event structure for a later venue.

2. Review of VerbNet

VerbNet is a lexicon of around 5,200 English verbs, organized primarily around Levin's (1993) verb classification. Classes in VerbNet are structured according to the verb's syntactic behavior. As described in (Kipper et al., 2006; Palmer, 2009; Bonial et al., 2011), VerbNet describes the sets of diathesis alternations that are compatible with each verb in the lexicon. For example, the verb *break* expresses both an inchoative form as well as a causative form, as already encountered in (1) above. Verbs such as *appear*, however, are compatible with an inchoative form (*A cloud appeared.*), but not in a causative construction. Classes are arranged hierarchically, with subclasses of verbs inheriting all the characteristics and frames of the parent class but exhibiting additional syntactic alternations.

Although the basis of the classification is largely syntactic, the verbs of a given class do share semantic regularities as well because, as Levin hypothesized, the syntactic behavior of a verb is largely determined by its meaning. Each class contains semantic predicates that are compatible with the member verbs and the class's syntactic frames. The semantic representations describe the participants at various stages of the event. For example, the representations for the **break** class, which includes such verbs as *shatter*, *snap*, and *tear*, describe a general Initial_state at the start of the event and a general Result at the end of the event.

(6) a. **break**: [NP V NP]
b. **example**: "Tony broke the window."
c. **syntax**: Agent V Patient
d. **semantics**: path_rel(start(E), Initial_State, Patient, change_of_state) & path_rel(end(E), Result, Patient, change_of_state)& cause(Agent, E) & contact(during(E), Instrument, Patient) & degradation_material_integrity(result(E), Patient)& physical_form(result(E), form, Patient)

The class does not refer to the type of contact that occurs or the specific form that results, although such distinctions could be made for subgroups of the class's verbs.

The related class **calibratable change of state**, covers events of change along a scale, such as *rise*, *fluctuate*, and *dwindle*. Its semantic representation makes no mention of contact or a degradation of material integrity. However, it also uses the path_rel *start(E)* and *end(E)* predicates, but substitutes change_on_scale for change_of_state and adds the predicate change_value(during(E), Patient, Direction). The direction is left underspecified, and no reference is made to any manner of the change, such as its speed.

3. VerbNet Predicative Structure

The first proposed change to VerbNet's semantic representation involves an enrichment to the predicative content associated with subevents that will help differentiate the meaning of within-class verbs. We believe that GL provides a framework with which to perform this kind of semantic componential analysis of word classes. To this end, there are two aspects of GL's semantic structure that will prove useful: predicate opposition structure and subeventual componential analysis. In addition, recent work on scalarity provides useful insights into how to distinguish verb classes involving incremental change (Kennedy and Levin, 2008).

Without an explicit representation of change of state, the lexical structure for a verb does not adequately model change dynamically. For this reason, the concept of *opposition structure* was introduced in GL as an enrichment to event structure (Pustejovsky, 2000). This makes explicit which predicate opposition is lexically encoded in a verb. For example, the verbs *die* and *kill* are both encoded with the opposition structure $[\neg dead(x), dead(x)]$. A binary opposition such as this can have distinct grammatical consequences, and this is reflected in VerbNet by membership in a specific class of change_of_state (COS), i.e., class 45.[4] In fact, identification of the mode of change and the scale associated with that change goes a long way towards explaining much of the grammatical behavior of such verbs (Hay et al., 1999; Kennedy and Levin, 2008).

VerbNet classes are motivated on the basis of syntactic and alternation-based behavior. We believe that it is possible to also identify semantically coherent clusters of verbs within these classes. A few examples will suggest our approach. Using GL-inspired componential analysis applied to the *run*-class (Verbnet 51.3.2), six distinct semantic dimensions emerge, which provide clear differentiations in meaning within this class. They are: 1) SPEED: *amble, bolt, sprint, streak, tear, chunter, flit, zoom*; PATH SHAPE: *cavort, hopscotch, meander, seesaw, slither, swerve, zigzag*; PURPOSE: *creep, pounce*; BODILY MANNER: *amble, ambulate, backpack, clump, clamber, shuffle*; ATTITUDE: *frolic, lumber, lurch, gallivant*; ORIENTATION: *slither, crawl, walk, backpack*. The benefit that such component-based analysis provides, as pointed out above, is that within-class semantic distinctions can be identified and also associated with behavior.

[4]The verb *die* is not formally marked as COS in VerbNet, but we can ignore this for the present discussion.

Theoretically inspired distinctions in meaning (e.g., the two motion verb classes of *path* and *manner*), can be systematically associated with (and hence identified with) specific grammatical realizations in the language. That is, given the right semantic vocabulary, linking components to syntactic behavior in the language can be annotated and then used for training classifiers and clustering algorithms. What is interesting about the class distinctions above, is that each dimension links to (associates with) clusters of syntactic clues and constructions. For example, PURPOSE associates with rationale and purpose clauses; the SPEED, ORIENTATION, and ATTITUDE dimensions select for adverbials for those attributes, respectively. This is the underlying benefit of deep semantic modeling: revealing underlying aspects of the event that are expressed syntactically, given a rich enough description, and an annotation strategy over datasets.

Let us return briefly to the examples mentioned in Section 1.0., where the verbs *slide*, *float*, and *roll* are all annotated as the same class, slide-11.2. We wish to identify those predicative forms that will sufficiently distinguish the meanings of these verbs. As pointed out in (Mani and Pustejovsky, 2012), the manner introduced by a verb such as *slide* is a *mereo-topological* specialization in meaning of a generic directed motion verb. This means that the nature of the movement is definable in terms referring to spatial configurations between an object (the ground) and the mover — or part of the mover (the figure). For example, *slide* and *roll* presuppose different modes of contact of the figure's surface to the ground, as well as presupposing a component of rotational symmetry for the figure. The VerbNet entries should reflect this distinction, which will entail reference to a Ground (G) role that is not currently part of the role inventory. Assuming such a participant (or something similar) is added to the inventory of roles in VerbNet, we can introduce the relation of "contact" between the mover and the ground to account for the first distinction, and a predicate for "rotational symmetry" to distinguish the second.

Using these features, we can distinguish several of Levin's classes of manner (including the members of **slide-11.2-1**), where a class is defined by certain constraints that hold throughout the event, E. For designating contact, we adopt RCC8's relations of "externally connected" (EC) and "disconnected" (DC) (Randell et al., 1992). To account for rotational symmetry, we introduce a relation between the moving object and its surface, which is in contact with the ground, i.e.,"rot-surface". These predicates facilitate three basic distinctions within this class: whether the mover is in touching the relative ground (*slide* vs. *fly*), when it it touching it (*slide* vs. *bounce*, and how it is touching it (*slide* vs. *roll*). Consider the definitions in (7).

(7) **Mereo-topological Distinctions**:
For Figure (F) relative to Ground (G):
a. EC(F,G), throughout E:
b. DC(F,G), throughout E:
c. (EC(F',G), throughout E, where rot-surface(F',F):
d. (EC(F,G), DC(F,G))*, throughout E.

For example, (7d) expresses the iterating step-wise motion involved in bouncing or hopping, where contact is followed by no contact, iterated throughout the event. That in (7c) expresses the condition present for a rotating surface in contact with the ground, i.e., *roll*. Finally, (7a) holds for motion of an object, F, involving continuous contact with the surface of the ground, G, while (7b) holds for motion with no contact between F and G. This distinguishes the verbs *slide* and *roll* from *float* and *fly*. The VerbNet representations with these distinction, for *slide* and *roll* might look like the following:

(8) a. **[NP V]**: motion(during(E), Figure) & while(E, EC(Figure,Ground))
b. **[NP V]**: motion(during(E), Figure) & while(E, EC(F',Ground)) & rot-surface(F',Figure)

This helps clarify the distinction between continuous contact verbs, such as *roll*, *drive*, and *walk*, from *float* and *fly*. This also has consequences when these verb classes each compose with orientational prepositions such as *over*, as illustrated in (9).

(9) a. The ball rolled over the grass.
(contact with the grass)
b. The balloon floated over the grass.
(no contact with the grass)

This illustrates that, while the orientation introduced by *over* is preserved in both classes, the semantics of contact is conveyed by the motion verb itself.

Finally, consider briefly the distinctions in VerbNet between the change_of_state classes, two of which were discussion in Section 2 above.

(10) a. 45.1: break-45.1
b. 45.2: bend-45.2
c. 45.3: cooking-45.3
d. 45.4: other_cos-45.4
e. 45.5: entity_specific_cos-45.5
f. 45.6.1: calibratable_cos-45.6.1
g. 45.6.2: caused_calibratable_cos-45.6.2
h. 45.7: remedy-45.7
i. 45.8: break_down-45.8

For each class, we propose that the opposition structure be explicitly encoded. Further, the nature of the scale structure should be identified, differentiating the following: what scale theory is assumed (nominal, binary, ordinal, interval, ratio); the attribute undergoing change; and whether the predicate denoting the attribute is associated with an *open* or *closed* scale. Through a similar strategy of differential semantic analysis applied across these classes, the nature of the change can be characterized using the vocabulary of GL qualia structure and types. For example, 45.1 involves an opposition structure over the FORMAL qualia role (denoting material integrity), while 45.2 refers to an aspect of the FORMAL, i.e., its "shape". The calibratable change verbs of 45.6.1 are incremental change predicates that are identified as changing along a specific attribute, whether the scale is open or closed, and the nature of the scale theory.

4. Conclusion

In this brief note, we have reported on some aspects of a new research project aimed at enriching VerbNet's predicative representations. This involved the introduction of systematic predicative enrichment to the verb's predicate structure. One part of this is an explicit identification of the mode of opposition structure inherent in the predicate. Another strategy involved GL-inspired semantic componential analysis over VerbNet classes.

We are also currently investigating a second modification concerning VerbNet's event representation, where we are studying how to integrate aspects of the event structure from GL (Pustejovsky, 1995a), specifically the notion of Dynamic Event Models (Pustejovsky and Moszkowicz, 2011; Pustejovsky, 2013) and Dynamic Argument Structure (Jezek and Pustejovsky, 2016). This is a significant issue, since VerbNet aims to represent the subeventual properties of the event as it unfolds, and it is important to ensure that the representation is both systematic and compositional in nature. This is a topic for ongoing research within this effort.

Acknowledgements

This work was supported to Brandeis by Contract W911NF-15-C-0238 with the US Defense Advanced Research Projects Agency (DARPA), to University of Colorado by Subaward 2015-06166-02 from UIUC for DARPA, and to both Brandeis and CU by the Defense Threat Reduction Agency (DTRA). Approved for Public Release, Distribution Unlimited. Thanks to comments from Elisabetta Jezek and Nikhil Krishnaswamy. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. All errors and mistakes are, of course, the responsibilities of the authors.

Claire Bonial, William Corvey, Martha Palmer, Volha V Petukhova, and Harry Bunt. 2011. A hierarchical unification of lirics and verbnet semantic roles. In *Semantic Computing (ICSC), 2011 Fifth IEEE International Conference on*, pages 483–489. IEEE.

Susan Windisch Brown, Dmitriy Dligach, and Martha Palmer. 2014. Verbnet class assignment as a wsd task. In *Computing Meaning*, pages 203–216. Springer.

William Croft. 2001. *Radical construction grammar: Syntactic theory in typological perspective*. Oxford University Press.

Adele E Goldberg and Ray Jackendoff. 2004. The english resultative as a family of constructions. *Language*, pages 532–568.

Jennifer Hay, Christopher Kennedy, and Beth Levin. 1999. Scalar structure underlies telicity in" degree achievements". In *Semantics and linguistic theory*, volume 9, pages 127–144.

Jena D Hwang, Annie Zaenen, and Martha Palmer. 2014. Criteria for identifying and annotating caused motion constructions in corpus data. In *LREC*, pages 1297–1304.

Elisabetta Jezek and James Pustejovsky. 2016. Dynamic argument structure.

Christopher Kennedy and Beth Levin. 2008. Measure of change: The adjectival core of degree achievements. *Adjectives and adverbs: Syntax, semantics and discourse*, pages 156–182.

Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2006. Extending verbnet with novel verb classes. In *Proceedings of LREC*, volume 2006, page 1. Citeseer.

Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.

Inderjeet Mani and James Pustejovsky. 2012. *Interpreting Motion: Grounded Representations for Spatial Language*. Oxford University Press.

Martha Palmer. 2009. Semlink: Linking propbank, verbnet and framenet. In *Proceedings of the Generative Lexicon Conference*, pages 9–15.

James Pustejovsky and Federica Busa. 1995. Unaccusativity and event composition. *Temporal reference, aspect, and actionality*, 1:159–177.

James Pustejovsky and Jessica Moszkowicz. 2011. The qualitative spatial dynamics of motion. *The Journal of Spatial Cognition and Computation*.

J. Pustejovsky. 1995a. *The Generative Lexicon*. Bradford Book. Mit Press.

James Pustejovsky. 1995b. *The Generative Lexicon*. MIT Press, Cambridge, MA.

J. Pustejovsky. 2000. Events and the semantics of opposition. In C. Tenny and J. Pustejovsky, editors, *Events as Grammatical Objects*, pages 445–482. Center for the Study of Language and Information (CSLI), Stanford, CA.

James Pustejovsky. 2013. Dynamic event structure and habitat theory. In *Proceedings of the 6th International Conference on Generative Approaches to the Lexicon (GL2013)*, pages 1–10. ACL.

D.A. Randell, Z. Cui, A. Cohn, B. Nebel, C. Rich, and W. Swartout. 1992. A spatial logic based on regions and connection. In *KR'92. Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*, pages 165–176, San Mateo. Morgan Kaufmann.

Chapter 5

Purpose Verbs. *Christiane Fellbaum*

Chapter 16

PURPOSE VERBS

Christiane Fellbaum
Princeton University, and Berlin-Brandenburgische Akademie der Wissenschaften

Abstract: Analogously to the distinction between TYPE and ROLE nouns, we differentiate between MANNER and PURPOSE verbs. Purpose verbs like exercise, treat, and cheat can conflate with manner verbs and contribute an additional, telic meaning component to these verbs. Conflation is triggered by contextual factors that create an “expectation” favoring a purpose interpretation over a pure manner reading. We compare and contrast purpose verbs with Functional Events in the Generative Lexicon. Some Functional Events are also purpose verbs, but Functional Events comprise a much larger and more loosely defined class.

Key words: lexical semantics, events, manner verbs, purpose verbs, functional events.

1. INTRODUCTION

Work on ontologies and lexical semantics has long been aware of different subsumption relations among concepts that are lexically encoded as nouns (Gangemi et al. 2001, Gangemi et al. 2002, Guarino and Welty 2001). An important distinction is that between TYPES and ROLES (Guarino and Welty 2002, Pustejovsky 1995). For example, poodle and Welsh corgi are TYPES of dog, while pet, hunting dog, guard dog, and working dog are ROLES associated with dogs. This distinction has important consequences for the structure of an ontology, its potential for reasoning, and its usefulness in AI and NLP applications.

The Generative Lexicon observes the difference between types and roles, labeling the latter Functional Types. Functional Types have a telic role associated with them (perhaps companionship for pet) and an intentionality that is introduced by an Agent.

While this distinction among entities, lexicalized as nouns, is fairly established, little work has been done on events, lexicalized as verbs. The Generative Lexicon distinguishes Functional Events, which are characterized by an telic and/or agentive role in the qualia structure of their arguments.

We examine the distinction between the two different event types from the perspective of WordNet, where verbs like run and speak are distinguished from verbs like exercise and greet. We argue for a classification into "manner" and "purpose" verbs; purpose verbs overlap only partially with the Functional Events of the Generative Lexicon.

Section 2 discussed the polysemy of the manner relation as it is coded in WordNet (Miller 1995, Fellbaum 1998). In Section 3, we argue that manner is far too broad a label, hiding in fact at least two distinct relations. This section motivates the distinction drawn in Section 4, which introduces purpose verbs and examines their property as distinct from other, established verb classes. In section 5, purpose verbs are compared to Functional Events in the Generative Lexicon. Section 6 is concerned with the representation of manner vs. purpose verbs, and section 7 discusses the distribution of purpose verbs in the lexicon.

2. VERBS IN WORDNET

WordNet's approach to the structure of the lexicon is to view it as a large network where each word is linked via one or more semantic relations to other words. The most important relation among linking verbs in WordNet is the manner relation (Fellbaum 1990, 1998).

MANNER is frequently taken to be a semantic primitive that defies further analysis (Wierzbicka 1996). In the lexical-conceptual structures of many verbs a MANNER component is assumed whose presence may have syntactic consequences (Rappaport, Hovav and Levin 1998, Hale and Keyser 1993, Krifka 1999, Jackendoff 1990, Talmy 1985), inter alia. Yet the exact nature of MANNER has never been made explicit.

At the same time, MANNER clearly plays an important role in verb meaning and structuring the lexicon. The WordNet experiment has shown that an intuitive notion of MANNER allows one to distinguish verbs and arrange them into tree-like hierarchies, with verbs denoting events that are increasingly semantically specified as one descends the hierarchy (Fellbaum 1990, 1998). WordNet makes use of a MANNER relation that constitutes a kind of counterpart to the ISA relation among nouns in WordNet and to subtyping in the Generative Lexicon.

One verb can be said to be subordinate of another verb when it denotes an event with an additional manner component that is missing in the less elaborate superordinate (Levin and Rapoport 1988). For example, stammer, lisp, and whisper are among the many manner subordinates of speak, as the statement “to stammer/lisp/whisper is to speak in some manner” shows. Manners of walking include ambling, slouching; splinter, crumble, and crush are among the verbs elaborating specific manners of break. And so forth.

Similarly to subtyping in the noun lexicon, it turns out that the manner relation is quite well suited to relate verb meanings to one another. WordNet has over 13,500 verb synonym sets; the vast majority are manner elaborations of some 500 basic verbs.

(Fellbaum 1998) points out that MANNER, as it is used in WordNet's hierarchical structures, is highly underspecified. Depending on the semantic domain, the differentiae distinguishing a base verb and a more elaborate subordinate may be dimensions like SPEED (walk-run), DIRECTION (move-rise), VOLUME (talk-scream), INTENSITY (persuade-brainwash), etc.

But (Fellbaum 2002a, 2002b) noted that WordNet's verb hierarchies ignore a more fundamental distinction among the concepts expressed by verbs.

3. TWO TYPES OF MANNER RELATIONS

Two apparently different relations can be found among verbs and their semantically elaborated manner subordinates. The distinction between the relations reveals a difference among types of verbs and the associated concepts, and parallels the distinction between type and roles in the noun lexicon drawn in the Generative Lexicon.

Consider the verb exercise on the one hand and verbs like jog, swim, and bike on the other hand. Jog, swim, and bike refer to manners of exercising, but they are clearly also manners of moving/travelling . Both the following statements are true:

(1) to jog/swim/bike is to exercise in some manner

(2) to jog/swim/bike is to move in some manner

But clearly, there is a difference. The relation between jog, swim, bike and exercise is defeasible: Not every jogging/swimming/biking event is necessarily an exercising event. By contrast, every jogging/swimming/biking event is necessarily a moving event:

(3) She jogged/swam/biked but did not exercise

(4) *She jogged/swam/biked but did not move

The concept exercise is definable only by means of subordinates like swim, jog, and bike that are shared with another subordinate, move. But move has many subordinates that are not shared with exercise, such as fly and drive.

The relation of jog, swim and bike to their superordinates move and exercise is similar to that between, e.g., dog, cat, and goldfish to animal on the one hand and to pet on the other hand:

(5) A dog/cat/goldfish is a kind of pet.

(6) A dog/cat/goldfish is a kind of animal.

(7) That's a dog/cat/goldfish, but it is not a pet.

(8) *That's a dog/cat/goldfish, but it is not an animal.

Just as one can recognize dogs, cats, and goldfish as animals, but not (necessarily) as pets (Guarino 1998), so one can recognize instances of biking, swimming, jogging as moving events, but not (necessarily) as exercising events. Unlike moving, the exercise component of biking, swimming, and jogging does not supply an identity criterion and is notionally dependent. Moving, but not exercising, is a necessary component of a biking/swimming/jogging event. So verbs like exercise seem similar to role nouns like pet, and verbs like move seems similar to type nouns like animal.

A random search in WordNet shows up a fair number of defeasible subsumption cases. One example is treat. A medical practitioner can treat a patient by massaging, injecting, bleeding, etc. But none of these necessarily constitute a treatment. A statement like “massaging (someone) is a manner or treating (him)” is not necessarily true, whereas the statement “massaging (someone) is manually manipulating (his body)” is necessarily true. So massaging is necessarily a manner of manipulating, but not necessarily a manner of treating.

4. PURPOSE VERBS

What kind of concepts are encoded by verbs like exercise, control, help, and treat, which may be, but are not necessarily, part of the network of verbs that can be constructed around MANNER? Unlike the non-defeasible superordinates of verbs like swim and massage, verbs like exercise etc. do not contribute a MANNER component to the meanings of their subordinate verbs. Instead, such verbs seem to express concepts that encode a kind of telicity or goal or purpose: One exercises, helps, treats, cheats, etc. with some goal or purpose in mind. This goal or purpose is generally intended by the agent.

We will refer to verbs like exercise, treat, cheat, control and help as PURPOSE VERBS, and we assume that their lexical-semantic structure includes a meaning component that could be labeled PURPOSE.

4.1 PURPOSE, MANNER, AND CHANGE-OF-STATE VERBS

A common distinction among verb classes is that between manner and change-of-state (COS) verbs (Rappaport, Hovav and Levin 1998), inter alia. We propose that purpose verbs constitute a third, distinct class.

4.2 PURPOSE AND MANNER

(Rappaport, Hovav and Levin 1998) observe that English verbs encode either a RESULT or a MANNER, but not both.

Similarly, we could not identify verbs that encode both a MANNER and a PURPOSE component as necessary parts of their lexical make-up, although we saw that manner verbs can be subordinates of purpose verbs in appropriate contexts.

Manner verbs do not say anything about a result that may ensue from the activity denoted by the verb. Resultant endstates may be encoded by secondary predicates:

(9) Tim wiped the table clean

(10) Kim shouted herself hoarse

(11) The couples waltzed themselves tired

By contrast, purpose verbs do not admit resultatives, even though many denote activities, an aspectual class that in principle admits resultatives:

(12) *I exercised myself strong

(13) *The doctor treated me healthy

(14) *The company cheated their stockholders poor

(15) *Paul helped Sue safe

(16) *The police controlled the crowd frightened

The fact that purpose verbs do not pattern with manner verbs further indicates that they do not contain a MANNER component.

4.3 PURPOSE AND CHANGE-OF-STATE

COS verbs, like purpose verbs, do not refer to MANNER. There are many ways of breaking a vase or of opening a door, and the manner in which a COS was effected may be stated in an adjunct phrase. Because the resultant state is expressed in the verb, no further resultative phrase is admitted:

(17) *Tim destroyed the painting ruined

(18) *Kim shredded the document illegible

Note that PP resultatives may be admissible, as in Kim shredded the documents into small pieces. (Fong, Fellbaum and Lebeaux 2001) distinguish several types of resultatives and their compatibility with different verb classes. Following their distinction, a verb like shred denotes a TRANSFORMATION, rather than a COS.

But purpose verbs are distinct from COS verbs in some important ways. First of all, purpose verbs may be activities, whereas COS verbs are always accomplishment or achievements, as the standard tests (Vendler 1967) show:

(19) He exercised for hours

(20) *He exercised in two hours

(21) The doctor treated me for years with the wrong medicine

(22) *The doctor treated me in minutes with the wrong medicine

(23) *He broke the vase for hours

(24) He broke the vase in seconds

(25) *She shredded the letter for days

(26) She shredded the letter in minutes

Second, COS verbs are causatives and have corresponding intransitives:

(27) Tim opened the door

(28) The door opened

(29) Kim broke the vase

(30) The vase broke

Transitive purpose verbs do not share this syntactic alternation:

(31) The doctor treated the patient

(32) *The patient treated

(33) The police controlled the crowd

(34) *The crowd controlled

nor are intransitive purpose verbs unaccusatives, as their aspectual properties show, as in (19).

However, transitive purpose verbs freely enter into middle constructions:

(35) The lawn mower controls easily

(36) Naive customers cheat easily

The subject in middle verbs is commonly referred to as “affected” (Keyser and Roeper 1984, Fellbaum 1985, Fagan 1988), inter alia. Affectedness is commonly treated as an unanalyzable primitive and has not received a precise semantic characterization. In particular, it is unclear how it contrasts with the notion “change of state.” Whatever the exact semantics of these concepts may be, purpose and COS verbs indicate that that there is a real difference between them. While COS verbs change the state of the Theme, purpose verbs merely affect them.

4.4 PURPOSE VERBS AND ADVERBS

The PURPOSE component of verbs like exercise and treat has an effect on the selection and interpretation of adverbs that co-occur with these verbs.

(Pustejovsky 1995) offers a Generative Lexicon account for the polysemy of adjectives. He notes that the telic role of nouns binds selectively with adjectives modifying the nouns, and that this process accounts for the appropriate reading of polysemous adjectives. For example, a fast car is a car that drives fast: the adjective is interpreted with respect to the telic role of the noun, which is drive. A different reading of the adjective obtains in the phrase fast typist, where the adjective is interpreted with respect to the telic role of typist, namely type.

Similarly, the purpose component of a verb appears to interact with certain adverbial modifiers. First, only purpose verbs select adverbs like (un)successfully, (in)effectively, fruitlessly, and with(out) result that modify the outcome of the event:

(37) John exercised with good results.

(38) Peter cheated successfully.

(39) Mary treated the patient effectively.

Such adverbs cannot be interpreted with manner verbs whose meanings lack a goal or purpose:

(40) ? John limped (un)successfully.

(41) ? Mary murmured fruitlessly.

(42) ? Kim scribbled effectively

Second, polysemous verbs with distinct manner and purpose readings are disambiguated by adverbs like (un)successfully. In the examples below, two different readings of run are accessed. (43) refers to a motion event, modified by a manner adverb. The event in the second sentence is interpreted as a competition or political race; run here is easily assigned the meaning run for office, i.e., a purpose verb.

(43) John ran fast

(44) John ran (un)successfully

The adverb may force a verb reading that assumes a purpose or goal:

(45) John spoke successfully

Although speak does not have an inherent purpose, (45) can be interpreted as a speaking event for a political purpose or a debate. Such a reading appears impossible for verbs that have a strong manner component:

(46) ? John stammered/stumbled/limped successfully

While verbs like like speak and run can be coerced into a purpose verb reading reading in the presence of adverbs like successfully, COS verbs cannot receive a purpose reading even when modified by such adverbs:

(47) ? She opened the door with good results

(48) ? He cracked the box successfully

5. FUNCTIONAL EVENTS IN THE GENERATIVE LEXICON

The Generative Lexicon (Pustejovsky 1995, 2001) classifies some verbs as Functional Events. Examples given in (Pustejovsky 2001) are eat, feed, run, greet, and spoil. Functional events are characterized by telic and/or agentive roles in the qualia of the verbs' arguments, i.e., the semantics of the verb arise from those of its arguments. Eating, feeding, and running are classified as Functional Events because they require agentivity and intention (Asher and Pustejovsky 2000). Similarly, a statement such as the food spoiled can be made only by an entity capable of judging the spoiled food's state with respect to its telic role (presumably, “nourishment”).

The telic role here that defines the event as functional is that of the verb's argument (food). By contrast, the purpose or goal that defines a purpose verb resides in the event and not in the telic role of the arguments. For example, the purpose of a greeting event is to acknowledge someone's presence, show recognition or kindness, etc. A purpose or goal presupposes an agentivity and intention, but not the telic roles of the participants.

Functional Events are defined intuitively rather than rigorously in (Pustejovsky 2001). And intuitively, there is some overlap between Functional Events and purpose verbs. Functional Events include purpose verbs, but the broad definition of Functional Events further encompasses verbs that are not purpose verbs. Beyond the agentivity and intentionality for Functional Events like eat cited by (Pustejovsky 2001), purpose verbs imply a purpose or telicity of the event that the Agent has in mind. A Functional Event like eat does not clearly express such a purpose, although the Agent involved in an eating event acts intentionally. By contrast, a purpose verb like greet, which is also classified by (Pustejovsky 2001, Asher and Pustejovsky 2000) as a Functional Event, qualifies as a purpose verb under the distinction proposed here.

To clarify the distinction, recall that a purpose verb like greet does not encode a manner: one can greet someone by nodding, waving, or pronouncing a greeting formula. Rather, greet expresses the purpose of a nodding, waving, or speaking event.

Another difference is that the Generative Lexicon's Functional Events, such as eating and running, are always recognizable as such, independent of the situational context. But labeling an event with a purpose verb like greet may depend on a subjective interpretation of that event. A nodding or waving event is not necessarily a greeting event, while a running event will be recognized and labeled as such by every observer.

There is a further difference between verbs like eat and run on the one hand, and verbs like greet on the other hand, which indicates that including them all in the category of Functional Events is too broad. In the case of verbs like eat and run, their relation to more specified manner verbs like munch and jog is not defeasible:

(49) *She munches but does not eat

(50) *They jog but don't run

But the relation of manner-of-greeting verbs like nod of wave to the base verb greet is defeasible:

(51) His waving/nodding is not a greeting

The distinction amonge verbs like munch and jog on the one hand, and greet on the other hand, is erased in the Generative Lexicon, where all these verbs are subsumed under the category of Functional Events. We argue that munch and jog are manner verbs, distinct from purpose verbs like greet.

In conclusion, we argued that Functional Events, as characterized by (Pustejovsky 2001, Asher and Pustejovsky 2000), include purpose verbs as well as other verbs that are not purpose verbs. Like Functional Events, purpose verbs presuppose intention and volition and hence agentivity. But these are merely necessary, not sufficient, meaning components.

6. REPRESENTATION

How can one represent the distinct meanings of verbs with both “manner” and “purpose” readings in a semantic network like WordNet?

6.1 REGULAR POLYSEMY?

One possibility is to posit two senses for verbs like swim, bike and jog, each with a different superordinate, here move and exercise. Some traditional dictionaries take this route; for exampe, jog is represented in the American Heritage Dictionary as having distinct running and exercising senses. But this solution has two undesirable effects. One is that it increases polysemy and suggests, falsely, that the two readings are unrelated. Moreover, there are likely to be contexts allowing only for an underspecified reading.

More seriously, positing two distinct senses misses the fact that every instance of jogging-as-exercise is necessarily also an instance of moving.

One might ask whether the “manner/purpose” readings of verbs like jog reflect a kind of systematic polysemy that can be accounted for by means of productive rules, similar to those found in the noun lexicon (Apresyan 1973). However, we could find no patterns of manner/purpose polysemy in the verb lexicon. Moreover, verbs denoting events that can be manners of treating, controlling, or helping can be semantically heterogeneous and do not seem to admit of any regularity that can be captured by means of regular polysemy rules.

Instead, the readings of many verbs as events with a purpose appear to be construed in an ad-hoc fashion from context. We will examine this point in more detail later.

6.2 MULTIPLE INHERITANCE?

Verbs like jog and bike could be related via the same labelled MANNER pointer to two superordinate parent concepts, one link being necessary and another defeasible. However, the resultant “tangled hierarchy” is clearly unsatisfactory, as it implies that every jogging/swimming/biking event is both an exercising and a moving event, when in fact only the latter is true.

A better way to capture the relevant semantic facts is to introduce two distinct kinds of relation linking a single verb to two superordinate concepts. In addition to strict hyponymy, there would be a “parallel” hyponymy relation with the appropriate properties.

6.3 PARA-RELATIONS

(Cruse 1986), in discussing the TYPE-ROLE distinction among nouns, proposes a relation dubbed para-hyponymy for organizing nouns like dog and pet hierarchically. Like regular hyponymy, para-hyponymy admits the formula Xs and other Ys, where X is the subordinate and Y the superordinate: Both statements, dogs and other canines (type) and dogs and other pets (role) are good. This formula can easily be adopted for verbs, and fits both strict hyponymy and para-hyponymy:

(52) Biking/swimming/jogging and other manners of moving/travelling

(53) Biking/swimming/jogging and other manners of exercising

The but not-test for nouns (Cruse 1986) that shows defeasibility, can be readily applied to verbs:

(54) It's a walking/jogging/biking event but it's not an exercising event.

6.4 EXPECTATION

(Cruse 1986) characterizes para-hyponymy among nouns not in terms of logical necessity but “expectation.” Thus, there seems to be an “expectation” that a jogging event is an exercising event, even though jogging is not necessarily exercising. While intuitively convincing, the notion of “expectation” immediately raises several questions, in particular if one wants to co-opt it to represent verbs and the events they denote. How can expectation be characterized? Can it be quantified? How can verb pairs related by para-hyponymy be identified in the lexicon? And how do we know whether, say, a verb token jog in a text or utterance refers to an exercising event or (merely) to a running event?

To begin with, expectation often appears to be context-dependent rather than inherent in the concept. In some contexts, a given verb's interpretation as a para-hyponym is more salient, whereas in other context, its reading as a strict manner hyponym of another superordinate is more appropriate.

For example, the verb's interpretations as a manner of moving is more salient in (55-57), whereas in (58), the events are readily interpreted as exercise:

(55) The boat capsized and we had to swim to the shore.

(56) My car is in the repair shop so I'll bike to work.

(57) It started to rain heavily so she ran into the library.

(58) He swims/bikes/runs 3 miles every morning before work.

Some contexts seem to favor an underspecified reading:

(59) He jogged to the store.

Second, the degree of expectation may differ across verbs independently of specific contexts, but be part of their lexical make-up. For some verbs, the para-relation is stronger than the strict relation, and the reverse may be true for other verbs. For example, jog intuitively is more strongly associated with its defeasible superordinate exercise than with its logical superordinate run, move. This is reflected in the fact that some dictionaries have distinct running and exercising senses for jog, as noted earlier. Conversely, walk seems be more strongly associated with move that with exercise. Walk seems like a less canonical form of exercise than jog, and thus exhibits a weaker association with its defeasible hypernym and a correspondingly stronger link to its strict superordinate.

The relative frequency of one reading as compared to another presumably influences expectation. Just as, say, hawks as pets may be more conventional in certain cultures than in others, there are probably cultures where jogging and running are not done for exercise but, say, for pursuing game in a hunt.

7. PURPOSE VERBS AND PARA-HYPONYMY IN THE LEXICON

This paper has cited only a handful of examples for purpose verbs and para-relations. At this point, we don't know how many such verbs there are in the English lexicon. If we think of the lexicon as a structured ontology, e.g., a large semantic network, one might ask whether such verbs are distributed randomly or in a systematic fashion. Another open question is whether the kinds of concepts expressed by purpose verbs are universally lexicalized and to what extent.

Almost any verb that is a hyponym of move could be made a para-hyponym of exercise, just as a pet reading can be coerced for many animals. If one wants to code para-relations in a lexicon, it is important to avoid flooding it with links that reflect readings with very low expectancy. It would therefore be desirable to firm up intuitions about the relative strength or weakness of the (para)hyponymy relation with the aid of corpus data.

7.1 FINDING PARA-RELATIONS

(Fellbaum 2002b) discusses ways of finding cases of para-relations among verbs from corpus data, and cites examples of attested data found on the web by means of characteristic pattern searches. These patterns are frames such as

(60) ..and other ways of (Y-ing)

(61) to (X) is to (Y)

(62) to (X) is not to (Y)

These searches overgenerate, as the frames turn up cases of hyponymy involving both manner and purpose. Manual sorting leaves us with examples such as these:

(63) Befriending, listening and other ways of helping....
www.britishcouncil.org/sudan/science/

(64) Walking and other exercise use many muscles.
www.lungusa.org/diseases/exercise.html

(65) Swimming, running, biking, walking and other exercise that are at a time length of over 20 minutes..
www.pmssolutions.com/Hiddentruth.html

(66) ... shake hands, using the right hand, and explain that this is a way of greeting one another. Pair up children and allow them to practice shaking hands.
www.atozkidsstuff.com/math.html

(67) Tipping, leaving a gratuity, is a way of thanking people for their service.
www.istudentcity.com/stages/

These examples show that targeted corpus searches can reveal the semantic relations among verbs (see Hearst 1998 for a discussion of patterns to find other semantic relations).

8. SUMMARY AND CONCLUSIONS

We have identified a class of “purpose” verbs that includes exercise, cheat, help, and treat. Such verbs encode neither MANNER nor RESULT, but encode an event with telicity. Unlike in the case of Functional Event in the Generative Lexicon, the telicity of purpose verbs is inherent in the event rather than in an argument of the verb expressing that event. Purpose verbs differ in several other respects from Functional Events in the Generative Lexicon. Functional Events are rather intuitively defined, and no test for distinguishing them from other event types has been given. By contrast, purpose verb can be clearly distinguished from manner and COS verbs and are incompatible with secondary predicates expressing results. Purpose verbs can be clearly distinguished from manner verbs, as the relations to their superordinate is defeasible. Finally, we showed that certain adverbs bind with the purpose component of these verbs for an appropriate interpretation.

Several open questions remain. How many purpose verbs are there in the English verb lexicon, and where in the lexicon are they? Do purpose verbs follow specific lexicalization patterns, similar to manner verbs? (Fellbaum 2002a) discusses ways of collecting naturally attested cases of this relation from corpora.

Semantic relations that are not based on logical necessity but on expectations grounded in pragmatics or world knowledge are an important area for lexical and ontological research. But we need to know more about how expected readings are generated from contexts.

ACKNOWLEDGMENT

This work was supported in part by the Alexander-von-Humboldt Foundation through the Zukunftsinvestitionsprogramm and ARDA (Advanced Research and Development Agency). I am grateful to the editors of this volume and to the reviewers of earlier drafts for many valuable comments. Special thanks go to Benjamin Haskell, without whose help this paper would not have materialized.

REFERENCES

Apresyan, Y. (1973) Regular Polysemy. *Linguistics*, 147:5-32.

Asher, N. and J. Pustejovsky (2001) The Metaphysics of Words in Context. Manuscript, Brandeis University, Waltham, MA.

Cruse, A. (1986) *Lexical Semantics*. Cambridge: Cambridge University Press.

Fagan, S. (1988) The English Middle. *Linguistic Inquiry*, 19:181-203.

Fellbaum, C. (1985) Adverbs in Agentless Actives and Passives. *Papers from the Parasession on Causatives and Agentivity*, 21:21-31. Chicago, IL: University of Chicago Press.

Fellbaum, C. (1990) The English Verb Lexicon as a Semantic Net. *International Journal of Lexicography* 3/4:278-301.

Fellbaum, C. (ed.) (1998) *WordNet:An Electronic Lexical Database*. Cambridge, MA: MIT Press.

Fellbaum, C. (2002a) On the Semantics of Troponymy. In Green, R., S. Myang and C. Bean, (eds.) *Relations*, 23-34. Dordrecht: Kluwer.

Fellbaum, C. (2002b) Parallel Hierarchies in the Verb Lexicon. In Simov, K. (ed.) *Proceedings of the Ontolex02 Workshop on Ontologies and Lexical Knowledge Bases*, 27-31. Paris, France: ELRA.

Fellbaum, C. and J. Kegl (1989) Taxonomic Structure and Object Deletion in the English Verbal System. In DeJong, K. and Y. No (eds) *Proceedings of the Sixth Eastern States Conference on Linguistics*, 94-103. Columbus, OH: Ohio State University Press.

Fong, S., C. Fellbaum and D. Lebeaux (2001) Ghosts, Shadows, and Resultatives: The Lexical Representation of Verbs. *Traitement Automatique de Langues*, 42:755-784.

Gangemi, A., N. Guarino and A. Oltramari (2001) Conceptual Analysis of Lexical Taxonomies. *Proceedings of FOIS*, 285-296. Ogunquit, MA: ACM.

Gangemi, A., N. Guarino, A. Oltromari and S. Borgo (2002) Cleaning up WordNet's Top Level. In Singh, U. N. (ed) *Proceedings of the First Global WordNet Conference*, 109-121. Mysore, India: Center for Indian Languages.

Guarino, N. (1998) Some Ontological principles for Designing Upper Level Lexical Resources. *Proceedings of the First International Conference of Language Resources and Evaluation*, 527-534. Paris: ELRA.

Guarino, N., and C. Welty (2001) *Identity and Subsumption*. LADSEP-CNR Internal Report. Padua, Italy.

Guarino, N., and C. Welty (2002) Evaluating Ontological Decisions with Ontoclean. *Communications of the ACM*, 45/2:61-65.

Hale, K. and S. Keyser (1993) On Argument Structure and the Lexical Expression of Syntactic Relations. In Hale, K. and S. Keyser (eds.) *The View From Building Twenty*, 53-109. MIT Press.

Hearst, M. A. (1998) Automatic Discovery of WordNet Relations. In Fellbaum, C. (ed.) *WordNet: An Electronic Lexical Database*, 131-151. Cambridge, MA: MIT Press.

Jackendoff, R. (1990) *Semantic Structures*. Cambridge, MA: MIT Press.

Keyser, S. and T. Roeper (1984) On the Middle and Ergative Constructions in English. *Linguistic Inquiry* 15:381-416.

Krifka, M. (1999) Manner in Dative Alternation. *Proceedings of the 18th West Coast Conference on Formal Linguistics*. Cornell, NY: Cascadilla Press.

Levin, B. and T. Rapoport (1988) Lexical Subordination. *Proceedings of the Chicago Linguistic Society*, 275-289. Chicago, IL: University of Chicago Press.

Miller, G. A. (1995) WordNet: A Lexical Database for English. *Communications of the ACM*, 38.11:39-41.

Pustejovsky, J. (1995) *The Generative Lexicon*. Cambridge, MA: MIT Press.

Pustejovsky, J. (2001) Type Construction and the Logic of Concepts. In Bouillon, P. and F. Busa (eds.) *The Syntax of Word Meaning*. Cambridge: Cambridge University Press,

Rappaport Hovav, Malka, and Beth Levin (1998) Building Word Meanings. In Butt, M. and W. Geuder (eds.) *The Projection of Arguments: Lexical and Compositional Factors*, 97-174. Stanford, CA: CSLI Publications.

Talmy, L. (1985) Lexicalization Patterns: Semantic Structure in Lexical Forms. In Shopen, T. (ed.) *Language Typology and Semantic Description*, 57-149. Cambridge: Cambridge University Press.

Vendler, Z. (1967) *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.

Wierzbicka, A. (1996) *Semantics, Primes, and Universals*. Oxford, UK: Oxford University Press.

Chapter 6

Abstract Meaning Representation for Sembanking. *Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, Nathan Schneider*

Abstract Meaning Representation for Sembanking

Laura Banarescu
SDL
lbanarescu@sdl.com

Claire Bonial
Linguistics Dept.
Univ. Colorado
claire.bonial@colorado.edu

Shu Cai
ISI
USC
shucai@isi.edu

Madalina Georgescu
SDL
mgeorgescu@sdl.com

Kira Griffitt
LDC
kiragrif@ldc.upenn.edu

Ulf Hermjakob
ISI
USC
ulf@isi.edu

Kevin Knight
ISI
USC
knight@isi.edu

Philipp Koehn
School of Informatics
Univ. Edinburgh
pkoehn@inf.ed.ac.uk

Martha Palmer
Linguistics Dept.
Univ. Colorado
martha.palmer@colorado.edu

Nathan Schneider
LTI
CMU
nschneid@cs.cmu.edu

Abstract

We describe Abstract Meaning Representation (AMR), a semantic representation language in which we are writing down the meanings of thousands of English sentences. We hope that a sembank of simple, whole-sentence semantic structures will spur new work in statistical natural language understanding and generation, like the Penn Treebank encouraged work on statistical parsing. This paper gives an overview of AMR and tools associated with it.

1 Introduction

Syntactic treebanks have had tremendous impact on natural language processing. The Penn Treebank is a classic example—a simple, readable file of natural-language sentences paired with rooted, labeled syntactic trees. Researchers have exploited manually-built treebanks to build statistical parsers that improve in accuracy every year. This success is due in part to the fact that we have a single, whole-sentence parsing task, rather than separate tasks and evaluations for base noun identification, prepositional phrase attachment, trace recovery, verb-argument dependencies, etc. Those smaller tasks are naturally solved as a by-product of whole-sentence parsing, and in fact, solved better than when approached in isolation.

By contrast, semantic annotation today is balkanized. We have separate annotations for named entities, co-reference, semantic relations, discourse connectives, temporal entities, etc. Each annotation has its own associated evaluation, and training data is split across many resources. We lack a simple readable sembank of English sentences paired with their whole-sentence, logical meanings. We believe a sizable sembank will lead to new work in statistical natural language understanding (NLU), resulting in semantic parsers that are as ubiquitous as syntactic ones, and support natural language generation (NLG) by providing a logical semantic input.

Of course, when it comes to whole-sentence semantic representations, linguistic and philosophical work is extensive. We draw on this work to design an Abstract Meaning Representation (AMR) appropriate for sembanking. Our basic principles are:

- AMRs are rooted, labeled graphs that are easy for people to read, and easy for programs to traverse.
- AMR aims to abstract away from syntactic idiosyncrasies. We attempt to assign the same AMR to sentences that have the same basic meaning. For example, the sentences "he described her as a genius", "his description of her: genius", and "she was a genius, according to his description" are all assigned the same AMR.
- AMR makes extensive use of PropBank framesets (Kingsbury and Palmer, 2002; Palmer et al., 2005). For example, we represent a phrase like "bond investor" using the frame "invest-01", even though no verbs appear in the phrase.
- AMR is agnostic about how we might want to derive meanings from strings, or vice-versa. In translating sentences to AMR, we do not dictate a particular sequence of rule applica-

tions or provide alignments that reflect such rule sequences. This makes sembanking very fast, and it allows researchers to explore their own ideas about how strings are related to meanings.

- AMR is heavily biased towards English. It is not an Interlingua.

AMR is described in a 50-page annotation guideline.[1] In this paper, we give a high-level description of AMR, with examples, and we also provide pointers to software tools for evaluation and sembanking.

2 AMR Format

We write down AMRs as rooted, directed, edge-labeled, leaf-labeled graphs. This is a completely traditional format, equivalent to the simplest forms of feature structures (Shieber et al., 1986), conjunctions of logical triples, directed graphs, and PENMAN inputs (Matthiessen and Bateman, 1991). Figure 1 shows some of these views for the sentence "The boy wants to go". We use the graph notation for computer processing, and we adapt the PENMAN notation for human reading and writing.

3 AMR Content

In neo-Davidsonian fashion (Davidson, 1969), we introduce variables (or graph nodes) for entities, events, properties, and states. Leaves are labeled with concepts, so that "(b / boy)" refers to an instance (called b) of the concept boy. Relations link entities, so that "(d / die-01 :location (p / park))" means there was a death (d) in the park (p). When an entity plays multiple roles in a sentence, we employ re-entrancy in graph notation (nodes with multiple parents) or variable re-use in PENMAN notation.

AMR concepts are either English words ("boy"), PropBank framesets ("want-01"), or special keywords. Keywords include special entity types ("date-entity", "world-region", etc.), quantities ("monetary-quantity", "distance-quantity", etc.), and logical conjunctions ("and", etc).

AMR uses approximately 100 relations:

- **Frame arguments, following PropBank conventions.** :arg0, :arg1, :arg2, :arg3, :arg4, :arg5.

[1] AMR guideline: amr.isi.edu/language.html

LOGIC format:

∃ w, b, g:
instance(w, want-01) ∧ instance(g, go-01) ∧
instance(b, boy) ∧ arg0(w, b) ∧
arg1(w, g) ∧ arg0(g, b)

AMR format (based on PENMAN):

```
(w / want-01
  :arg0 (b / boy)
  :arg1 (g / go-01
          :arg0 b))
```

GRAPH format:

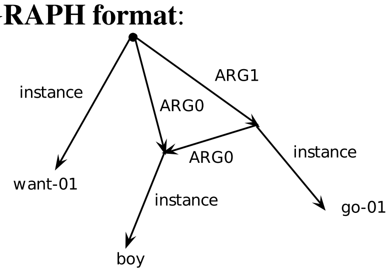

Figure 1: Equivalent formats for representating the meaning of "The boy wants to go".

- **General semantic relations.** :accompanier, :age, :beneficiary, :cause, :compared-to, :concession, :condition, :consist-of, :degree, :destination, :direction, :domain, :duration, :employed-by, :example, :extent, :frequency, :instrument, :li, :location, :manner, :medium, :mod, :mode, :name, :part, :path, :polarity, :poss, :purpose, :source, :subevent, :subset, :time, :topic, :value.

- **Relations for quantities.** :quant, :unit, :scale.

- **Relations for date-entities.** :day, :month, :year, :weekday, :time, :timezone, :quarter, :dayperiod, :season, :year2, :decade, :century, :calendar, :era.

- **Relations for lists.** :op1, :op2, :op3, :op4, :op5, :op6, :op7, :op8, :op9, :op10.

AMR also includes the inverses of all these relations, e.g., :arg0-of, :location-of, and :quant-of. In addition, every relation has an associated reification, which is what we use when we want to modify the relation itself. For example, the reification of :location is the concept "be-located-at-91".

Our set of concepts and relations is designed to allow us represent all sentences, taking all words into account, in a reasonably consistent manner. In the rest of this section, we give examples of how AMR represents various kinds of words, phrases, and sentences. For full documentation, the reader is referred to the AMR guidelines.

Frame arguments. We make heavy use of PropBank framesets to abstract away from English syntax. For example, the frameset "describe-01" has three pre-defined slots (:arg0 is the describer, :arg1 is the thing described, and :arg2 is what it is being described as).

```
(d / describe-01
  :arg0 (m / man)
  :arg1 (m2 / mission)
  :arg2 (d / disaster))

The man described the mission as a disaster.
The man's description of the mission:
  disaster.
As the man described it, the mission was a
  disaster.
```

Here, we do not annotate words like "as" or "it", considering them to be syntactic sugar.

General semantic relations. AMR also includes many non-core relations, such as :beneficiary, :time, and :destination.

```
(s / hum-02
  :arg0 (s2 / soldier)
  :beneficiary (g / girl)
  :time (w / walk-01
          :arg0 g
          :destination (t / town)))

The soldier hummed to the girl as she
  walked to town.
```

Co-reference. AMR abstracts away from co-reference gadgets like pronouns, zero-pronouns, reflexives, control structures, etc. Instead we re-use AMR variables, as with "g" above. AMR annotates sentences independent of context, so if a pronoun has no antecedent in the sentence, its nominative form is used, e.g., "(h / he)".

Inverse relations. We obtain rooted structures by using inverse relations like :arg0-of and :quant-of.

```
(s / sing-01
  :arg0 (b / boy
          :source (c / college)))

The boy from the college sang.

(b / boy
  :arg0-of (s / sing-01)
  :source (c / college))

the college boy who sang ...

(i / increase-01
  :arg1 (n / number
          :quant-of (p / panda)))

The number of pandas increased.
```

The top-level root of an AMR represents the focus of the sentence or phrase. Once we have selected the root concept for an entire AMR, there are no more focus considerations—everything else is driven strictly by semantic relations.

Modals and negation. AMR represents negation logically with :polarity, and it expresses modals with concepts.

```
(g / go-01
  :arg0 (b / boy)
  :polarity -)

The boy did not go.

(p / possible
  :domain (g / go-01
            :arg0 (b / boy))
  :polarity -))

The boy cannot go.
It's not possible for the boy to go.

(p / possible
  :domain (g / go-01
            :arg0 (b / boy)
            :polarity -))

It's possible for the boy not to go.

(p / obligate-01
  :arg2 (g / go-01
          :arg0 (b / boy))
  :polarity -)

The boy doesn't have to go.
The boy isn't obligated to go.
The boy need not go.

(p / obligate-01
  :arg2 (g / go-01
          :arg0 (b / boy)
          :polarity -))

The boy must not go.
```

```
It's obligatory that the boy not go.

(t / think-01
  :arg0 (b / boy)
  :arg1 (w / win-01
          :arg0 (t / team)
          :polarity -))

The boy doesn't think the team will win.
The boy thinks the team won't win.
```

Questions. AMR uses the concept "amr-unknown", in place, to indicate wh-questions.

```
(f / find-01
  :arg0 (g / girl)
  :arg1 (a / amr-unknown))

What did the girl find?

(f / find-01
  :arg0 (g / girl)
  :arg1 (b / boy)
  :location (a / amr-unknown))

Where did the girl find the boy?

(f / find-01
  :arg0 (g / girl)
  :arg1 (t / toy
          :poss (a / amr-unknown)))

Whose toy did the girl find?
```

Yes-no questions, imperatives, and embedded wh-clauses are treated separately with the AMR relation :mode.

Verbs. Nearly every English verb and verb-particle construction we have encountered has a corresponding PropBank frameset.

```
(l / look-05
  :arg0 (b / boy)
  :arg1 (a / answer))

The boy looked up the answer.
The boy looked the answer up.
```

AMR abstracts away from light-verb constructions.

```
(a / adjust-01
  :arg0 (g / girl)
  :arg1 (m / machine))

The girl adjusted the machine.
The girl made adjustments to the machine.
```

Nouns. We use PropBank verb framesets to represent many nouns as well.

```
(d / destroy-01
  :arg0 (b / boy)
  :arg1 (r / room))

the destruction of the room by the boy ...
the boy's destruction of the room ...
The boy destroyed the room.
```

We never say "destruction-01" in AMR. Some nominalizations refer to a whole event, while others refer to a role player in an event.

```
(s / see-01
  :arg0 (j / judge)
  :arg1 (e / explode-01))

The judge saw the explosion.

(r / read-01
  :arg0 (j / judge)
  :arg1 (t / thing
          :arg1-of (p / propose-01))

The judge read the proposal.

(t / thing
  :arg1-of (o / opine-01
             :arg0 (g / girl)))

the girl's opinion
the opinion of the girl
what the girl opined
```

Many "-er" nouns invoke PropBank framesets. This enables us to make use of slots defined for those framesets.

```
(p / person
  :arg0-of (i / invest-01))

investor

(p / person
  :arg0-of (i / invest-01
             :arg1 (b / bond)))

bond investor

(p / person
  :arg0-of (i / invest-01
             :manner (s / small)))

small investor

(w / work-01
  :arg0 (b / boy)
  :manner (h / hard))

the boy is a hard worker
the boy works hard
```

However, a treasurer is not someone who treasures, and a president is not (just) someone who presides.

Adjectives. Various adjectives invoke PropBank framesets.

```
(s / spy
  :arg0-of (a / attract-01))

the attractive spy

(s / spy
  :arg0-of (a / attract-01
             :arg1 (w / woman)))

the spy who is attractive to women
```

"-ed" adjectives frequently invoke verb framesets. For example, "acquainted with magic" maps to "acquaint-01". However, we are not restricted to framesets that can be reached through morphological simplification.

```
(f / fear-01
  :arg0 (s / soldier)
  :arg1 (b / battle-01))

The soldier was afraid of battle.
The soldier feared battle.
The soldier had a fear of battle.
```

For other adjectives, we have defined new framesets.

```
(r / responsible-41
  :arg1 (b / boy)
  :arg2 (w / work))

The boy is responsible for the work.
The boy has responsibility for the work.
```

While "the boy responsibles the work" is not good English, it is perfectly good Chinese. Similarly, we handle *tough*-constructions logically.

```
(t / tough
  :domain (p / please-01
            :arg1 (g / girl)))

Girls are tough to please.
It is tough to please girls.
Pleasing girls is tough.
```

"please-01" and "girl" are adjacent in the AMR, even if they are not adjacent in English. "-able" adjectives often invoke the AMR concept "possible", but not always (e.g., a "taxable fund" is actually a "taxed fund").

```
(s / sandwich
  :arg1-of (e / eat-01
             :domain-of (p / possible)))

an edible sandwich

(f / fund
  :arg1-of (t / tax-01))

a taxable fund
```

Pertainym adjectives are normalized to root form.

```
(b / bomb
  :mod (a / atom))

atom bomb
atomic bomb
```

Prepositions. Most prepositions simply signal semantic frame elements, and are themselves dropped from AMR.

```
(d / default-01
  :arg1 (n / nation)
  :time (d2 / date-entity
          :month 6))

The nation defaulted in June.
```

Time and location prepositions are kept if they carry additional information.

```
(d / default-01
  :arg1 (n / nation)
  :time (a / after
          :op1 (w / war-01)))

The nation defaulted after the war.
```

Occasionally, neither PropBank nor AMR has an appropriate relation, in which case we hold our nose and use a :prep-X relation.

```
(s / sue-01
  :arg1 (m / man)
  :prep-in (c / case))

The man was sued in the case.
```

Named entities. Any concept in AMR can be modified with a :name relation. However, AMR includes standardized forms for approximately 80 named-entity types, including person, country, sports-facility, etc.

```
(p / person
  :name (n / name
          :op1 "Mollie"
          :op2 "Brown"))

Mollie Brown

(p / person
```

```
  :name (n / name
          :op1 "Mollie"
          :op2 "Brown")
  :arg0-of (s / slay-01
             :arg1 (o / orc)))

the orc-slaying Mollie Brown
Mollie Brown, who slew orcs
```

AMR does not normalize multiple ways of referring to the same concept (e.g., "US" versus "United States"). It also avoids analyzing semantic relations inside a named entity—e.g., an organization named "Stop Malaria Now" does not invoke the "stop-01" frameset. AMR gives a clean, uniform treatment to titles, appositives, and other constructions.

```
(c / city
  :name (n / name
          :op1 "Zintan"))

Zintan
the city of Zintan

(p / president
  :name (n / name
          :op1 "Obama"))

President Obama
Obama, the president ...

(g / group
  :name (n / name
          :op1 "Elsevier"
          :op2 "N.V.")
  :mod (c / country
         :name (n2 / name
                 :op1 "Netherlands"))
  :arg0-of (p / publish-01))

Elsevier N.V., the Dutch publishing group...
Dutch publishing group Elsevier N.V. ...
```

Copula. Copulas use the :domain relation.

```
(w / white
  :domain (m / marble))

The marble is white.

(l / lawyer
  :domain (w / woman))

The woman is a lawyer.

(a / appropriate
  :domain (c / comment)
  :polarity -))

The comment is not appropriate.
The comment is inappropriate.
```

Reification. Sometimes we want to use an AMR relation as a first-class concept—to be able to modify it, for example. Every AMR relation has a corresponding reification for this purpose.

```
(m / marble
  :location (j / jar))

the marble in the jar ...

(b / be-located-at-91
  :arg1 (m / marble)
  :arg2 (j / jar)
  :polarity -)
  :time (y / yesterday))

The marble was not in the jar yesterday.
```

If we do not use the reification, we run into trouble.

```
(m / marble
  :location (j / jar
              :polarity -)
  :time (y / yesterday))

yesterday's marble in the non-jar ...
```

Some reifications are standard PropBank framesets (e.g., "cause-01" for :cause, or "age-01" for :age).

This ends the summary of AMR content. For lack of space, we omit descriptions of comparatives, superlatives, conjunction, possession, determiners, date entities, numbers, approximate numbers, discourse connectives, and other phenomena covered in the full AMR guidelines.

4 Limitations of AMR

AMR does not represent inflectional morphology for tense and number, and it omits articles. This speeds up the annotation process, and we do not have a nice semantic target representation for these phenomena. A lightweight syntactic-style representation could be layered in, via an automatic post-process.

AMR has no universal quantifier. Words like "all" modify their head concepts. AMR does not distinguish between real events and hypothetical, future, or imagined ones. For example, in "the boy wants to go", the instances of "want-01" and "go-01" have the same status, even though the "go-01" may or may not happen.

We represent “history teacher” nicely as “(p / person :arg0-of (t / teach-01 :arg1 (h / history)))”. However, “history professor” becomes “(p / professor :mod (h / history))”, because “profess-01” is not an appropriate frame. It would be reasonable in such cases to use a NomBank (Meyers et al., 2004) noun frame with appropriate slots.

5 Creating AMRs

We have developed a power editor for AMR, accessible by web interface.[2] The AMR Editor allows rapid, incremental AMR construction via text commands and graphical buttons. It includes online documentation of relations, quantities, reifications, etc., with full examples. Users log in, and the editor records AMR activity. The editor also provides significant guidance aimed at increasing annotator consistency. For example, users are warned about incorrect relations, disconnected AMRs, words that have PropBank frames, etc. Users can also search existing sembanks for phrases to see how they were handled in the past. The editor also allows side-by-side comparison of AMRs from different users, for training purposes.

In order to assess inter-annotator agreement (IAA), as well as automatic AMR parsing accuracy, we developed the *smatch* metric (Cai and Knight, 2013) and associated script.[3] Smatch reports the semantic overlap between two AMRs by viewing each AMR as a conjunction of logical triples (see Figure 1). Smatch computes precision, recall, and F-score of one AMR’s triples against the other’s. To match up variables from two input AMRs, smatch needs to execute a brief search, looking for the variable mapping that yields the highest F-score.

Smatch makes no reference to English strings or word indices, as we do not enforce any particular string-to-meaning derivation. Instead, we compare semantic representations directly, in the same way that the MT metric Bleu (Papineni et al., 2002) compares target strings without making reference to the source.

For an initial IAA study, and prior to adjusting the AMR Editor to encourage consistency, 4 expert AMR annotators annotated 100 newswire sentences and 80 web text sentences. They then created consensus AMRs through discussion. The average annotator vs. consensus IAA (smatch) was 0.83 for newswire and 0.79 for web text. When newly trained annotators doubly annotated 382 web text sentences, their annotator vs. annotator IAA was 0.71.

[2] AMR Editor: amr.isi.edu/editor.html

[3] Smatch: amr.isi.edu/evaluation.html

6 Current AMR Bank

We currently have a manually-constructed AMR bank of several thousand sentences, a subset of which can be freely downloaded,[4] the rest being distributed via the LDC catalog.

In initially developing AMR, the authors built consensus AMRs for:

- 225 short sentences for tutorial purposes
- 142 sentences of newswire (*)
- 100 sentences of web data (*)

Trained annotators at LDC then produced AMRs for:

- 1546 sentences from the novel “The Little Prince”
- 1328 sentences of web data
- 1110 sentences of web data (*)
- 926 sentences from Xinhua news (*)
- 214 sentences from CCTV broadcast conversation (*)

Collections marked with a star (*) are also in the OntoNotes corpus (Pradhan et al., 2007; Weischedel et al., 2011).

Using the AMR Editor, annotators are able to translate a full sentence into AMR in 7-10 minutes and postedit an AMR in 1-3 minutes.

7 Related Work

Researchers working on whole-sentence semantic parsing today typically use small, domain-specific sembanks like GeoQuery (Wong and Mooney, 2006). The need for larger, broad-coverage sembanks has sparked several projects, including the Groningen Meaning Bank (GMB) (Basile et al., 2012a), UCCA (Abend and Rappoport, 2013), the Semantic Treebank (ST) (Butler and Yoshimoto, 2012), the Prague Dependency Treebank (Böhmová et al., 2003), and UNL (Uchida et al., 1999; Uchida et al., 1996; Martins, 2012).

[4] amr.isi.edu/download.html

Concepts. Most systems use English words as concepts. AMR uses PropBank frames (e.g., "describe-01"), and UNL uses English WordNet synsets (e.g., "200752493").

Relations. GMB uses VerbNet roles (Schuler, 2005), and AMR uses frame-specific PropBank relations. UNL has a dedicated set of over 30 frequently used relations.

Formalism. GMB meanings are written in DRT (Kamp et al., 2011), exploiting full first-order logic. GMB and ST both include universal quantification.

Granularity. GMB and UCCA annotate short texts, so that the same entity can participate in events described in different sentences; other systems annotate individual sentences.

Entities. AMR uses 80 entity types, while GMB uses 7.

Manual versus automatic. AMR, UNL, and UCCA annotation is fully manual. GMB and ST produce meaning representations automatically, and these can be corrected by experts or crowds (Venhuizen et al., 2013).

Derivations. AMR and UNL remain agnostic about the relation between strings and their meanings, considering this a topic of open research. ST and GMB annotate words and phrases directly, recording derivations as (for example) Montague-style compositional semantic rules operating on CCG parses.

Top-down verus bottom-up. AMR annotators find it fast to construct meanings from the top down, starting with the main idea of the sentence (though the AMR Editor allows bottom-up construction). GMB and UCCA annotators work bottom-up.

Editors, guidelines, genres. These projects have graphical sembanking tools (e.g., Basile et al. (2012b)), annotation guidelines,[5] and sembanks that cover a wide range of genres, from news to fiction. UNL and AMR have both annotated many of the same sentences, providing the potential for direct comparison.

8 Future Work

Sembanking. Our main goal is to continue sembanking. We would like to employ a large sembank to create shared tasks for natural language understanding and generation. These tasks may additionally drive interest in theoretical frameworks for probabilistically mapping between graphs and strings (Quernheim and Knight, 2012b; Quernheim and Knight, 2012a; Chiang et al., 2013).

Applications. Just as syntactic parsing has found many unanticipated applications, we expect sembanks and statistical semantic processors to be used for many purposes. To get started, we are exploring the use of statistical NLU and NLG in a semantics-based machine translation (MT) system. In this system, we annotate bilingual Chinese/English data with AMR, then train components to map Chinese to AMR, and AMR to English. A prototype is described by Jones et al. (2012).

Disjunctive AMR. AMR aims to canonicalize multiple ways of saying the same thing. We plan to test how well we are doing by building AMRs on top of large, manually-constructed paraphrase networks from the HyTER project (Dreyer and Marcu, 2012). Rather than build individual AMRs for different paths through a network, we will construct highly-packed disjunctive AMRs. With this application in mind, we have developed a guideline[6] for disjunctive AMR. Here is an example:

```
(o / *OR*
  :op1 (t / talk-01)
  :op2 (m / meet-03)
  :OR (o2 / *OR*
        :mod (o3 / official)
        :arg1-of (s / sanction-01
                   :arg0 (s2 / state))))

official talks
state-sanctioned talks
meetings sanctioned by the state
```

AMR extensions. Finally, we would like to deepen the AMR language to include more relations (to replace :mod and :prep-X, for example), entity normalization (perhaps wikification), quantification, and temporal relations. Ultimately, we would like to also include a comprehensive set of more abstract frames like "Earthquake-01" (:magnitude, :epicenter, :casualties), "CriminalLawsuit-01" (:defendant, :crime, :jurisdiction), and "Pregnancy-01" (:father, :mother, :due-date). Projects like FrameNet (Baker et al., 1998) and CYC (Lenat, 1995) have long pursued such a set.

[5] UNL guidelines: www.undl.org/unlsys/unl/unl2005

[6] Disjunctive AMR guideline: amr.isi.edu/damr.1.0.pdf

9 Acknowledgements

This work was supported in part by NSF grant IIS-0908532.

References

O. Abend and A. Rappoport. 2013. UCCA: A semantics-based grammatical annotation scheme. In *Proc. IWCS*.

C. Baker, C. Fillmore, and J. Lowe. 1998. The Berkeley FrameNet project. In *Proc. COLING*.

V. Basile, J. Bos, K. Evang, and N. Venhuizen. 2012a. Developing a large semantically annotated corpus. In *Proc. LREC*.

V. Basile, J. Bos, K. Evang, and N. Venhuizen. 2012b. A platform for collaborative semantic annotation. In *Proc. EACL demonstrations*.

A. Böhmová, J. Hajič, E. Hajičová, and B. Hladká. 2003. The Prague dependency treebank. In *Treebanks*. Springer.

A. Butler and K. Yoshimoto. 2012. Banking meaning representations from treebanks. *Linguistic Issues in Language Technology*, 7.

S. Cai and K. Knight. 2013. Smatch: An accuracy metric for abstract meaning representations. In *Proc. ACL*.

D. Chiang, J. Andreas, D. Bauer, K. M. Hermann, B. Jones, and K. Knight. 2013. Parsing graphs with hyperedge replacement grammars. In *Proc. ACL*.

D. Davidson. 1969. The individuation of events. In N. Rescher, editor, *Essays in Honor of Carl G. Hempel*. D. Reidel, Dordrecht.

M. Dreyer and D. Marcu. 2012. Hyter: Meaning-equivalent semantics for translation evaluation. In *Proc. NAACL*.

B. Jones, J. Andreas, D. Bauer, K. M. Hermann, and K. Knight. 2012. Semantics-based machine translation with hyperedge replacement grammars. In *Proc. COLING*.

H. Kamp, J. Van Genabith, and U. Reyle. 2011. Discourse representation theory. In *Handbook of philosophical logic*, pages 125–394. Springer.

P. Kingsbury and M. Palmer. 2002. From TreeBank to PropBank. In *Proc. LREC*.

D. B. Lenat. 1995. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11).

R. Martins. 2012. Le Petit Prince in UNL. In *Proc. LREC*.

C. M. I. M. Matthiessen and J. A. Bateman. 1991. *Text Generation and Systemic-Functional Linguistics*. Pinter, London.

A. Meyers, R. Reeves, C. Macleod, R. Szekely, V. Zielinska, B. Young, and R. Grishman. 2004. The NomBank project: An interim report. In *HLT-NAACL 2004 workshop: Frontiers in corpus annotation*.

M. Palmer, D. Gildea, and P. Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1).

K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, Philadelphia, PA.

S. Pradhan, E. Hovy, M. Marcus, M. Palmer, L. Ramshaw, and R. Weischedel. 2007. Ontonotes: A unified relational semantic representation. *International Journal of Semantic Computing (IJSC)*, 1(4).

D. Quernheim and K. Knight. 2012a. DAGGER: A toolkit for automata on directed acyclic graphs. In *Proc. FSMNLP*.

D. Quernheim and K. Knight. 2012b. Towards probabilistic acceptors and transducers for feature structures. In *Proc. SSST Workshop*.

K. Schuler. 2005. *VerbNet: A broad-coverage, comprehensive verb lexicon*. Ph.D. thesis, University of Pennsylvania.

S. Shieber, F. C. N. Pereira, L. Karttunen, and M. Kay. 1986. Compilation of papers on unification-based grammar formalisms. Technical Report CSLI-86-48, Center for the Study of Language and Information, Stanford, California.

H. Uchida, M. Zhu, and T. Della Senta. 1996. UNL: Universal Networking Language—an electronic language for communication, understanding and collaboration. Technical report, IAS/UNU Tokyo.

H. Uchida, M. Zhu, and T. Della Senta. 1999. A gift for a millennium. Technical report, IAS/UNU Tokyo.

N. Venhuizen, V. Basile, K. Evang, and J. Bos. 2013. Gamification for word sense labeling. In *Proc. IWCS*.

R. Weischedel, E. Hovy, M. Marcus, M. Palmer, R. Belvin, S. Pradhan, L. Ramshaw, and N. Xue. 2011. OntoNotes: A large training corpus for enhanced processing. In J. Olive, C. Christianson, and J. McCary, editors, *Handbook of Natural Language Processing and Machine Translation*. Springer.

Y. W. Wong and R. J. Mooney. 2006. Learning for semantic parsing with statistical machine translation. In *Proc. HLT-NAACL*.