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.

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)

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

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 $\otimes$ ("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.

Argument is	Verb selects		
	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 predi-

cative 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

191

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, iro-
       ning, shopping
    b. E-I, Q-E of phys$\otimes_{telic}$ τ: penicillin, sandwich, cigarette, cake, des-
       sert, 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$\otimes_{\text{telic}} \tau$: 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$\otimes_{\text{telic}} \tau$: food
    c. E-I-Q-I of abstr$\otimes_{\text{telic}} \tau$: tax
    d. E-I-Q-I of location$\otimes_{\text{telic}} \tau$: 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$\otimes_{telic}$ τ: 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

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.