

A Distributional Model of Affordances in Semantic Type Coercion

Stephen McGregor
Laboratoire Lattice
CNRS & École normale supérieure / PSL
Université Sorbonne nouvelle Paris 3 / USPC
semcgregor@hotmail.com

Elisabetta Jezek
Theoretical and Applied Linguistics
Department of Humanities
University of Pavia
jezek@unipv.it

Abstract

We explore a novel application for interpreting semantic type coercions, motivated by insight into the role that perceptual affordances play in the selection of artefactual nouns that are observed as arguments for verbs that would stereotypically select for objects of a different type. In order to simulate affordances, which we take to be direct perceptions of context-specific opportunities for action, we perform a distributional analysis of dependency relationships between target words and their modifiers and adjuncts. We propose a novel methodology that uses these relationships as the basis for generating on-line transformations projecting semantic subspaces in which the interpretations of coercive compositions are expected to emerge as salient word-vectors. We offer some preliminary examples of how this model operates on phrases involving coercive verb-object interactions.

1 Introduction

As a linguistic phenomenon that is both elusive and pervasive, semantic type coercion presents computational models with a particular challenge. The problem is to find literal interpretations of coerced phrases which tend to seem quite natural to humans, and which are correspondingly observed at a high frequency in distributional analyses of large-scale corpora.

In what follows we present an overview of the phenomenon followed by a preliminary proposal for a context-sensitive framework for interpreting predicate-object coercions. Our methodology is inspired by theoretical insight into environmental affordances, and in this regard is in line with technical applications described in the area of image labelling by McGregor and Lim (2018). Motivated by an analysis of some of the shortcomings of a more general probabilistic approach, and also by a number of previous approaches to interpreting semantic coercion, we outline a model grounded in the distributional semantic modelling paradigm (Clark, 2015). In particular we propose a technique for constructing tensors based on an analysis of dependency relationships: this approach facilitates coercion interpretations (and conversely perhaps constructions) as geometric transformations, a move that might offer a plausible platform for capturing the direct perceptibility of environmental affordances (Raczaszek-Leonardi et al., 2018) using computational models. We present this work as an introductory overview, accompanied by a handful of examples, of a theoretically motivated methodology.

2 Background: Coercion, Affordances, and Distributional Semantics

Coercion is a theoretical tool that has been used in linguistic studies since Moens and Steedman (1988) to account for the fact that certain word combinations generate interpretations that are enriched or different from the strictly compositional ones. In the Generative Lexicon framework, predicate-argument coercion has been defined as the compositional mechanisms that resolves mismatches between the semantic type expected by a predicate for a specific argument position and the semantic type of the argument filler, by adjusting the type of the argument to satisfy the type requirement of the verb (*argument type coercion*; Pustejovsky, 1991). A classic example is (1), where *wine* is said to be coerced to an activity as a result of the semantic requirements the predicate imposes on its object, i.e. *finish* applies to an

activity.¹ The implicit activity is claimed to be “drinking”, and it is assumed to be stored in the lexical entry as a value of its telic *quale*.²

1 “When they finished the wine, he stood up”.

Another line of research in cognitive and distributional semantics proposes to frame coercion in terms of thematic role fit, defined as the semantic plausibility of an argument filler to fulfil the expectation of a verb in a given role, expressed in terms of score. Thematic fit scores range from 0 to 1, and correspond to the cosine similarity to the centroid, or vector average, computed over the most typical role fillers for that verb (Zarcone et al. (2013), Greenberg et al. (2015)). Under this view, coercion is interpreted as the result of low thematic fit scores of the fillers of the argument positions of a verb rather than the response to a type clash. For example, entity-denoting objects like *wine* have a low thematic fit as objects of event-selecting verbs like *finish*, and the recovery of the implicit event is seen as a consequence of the dispreference of the verb for the entity-denoting argument. In the thematic fit approach, the retrieval of the covert event relies on general event knowledge (GEK) activation (Chersoni et al., 2017).

In our proposal, we examine the phenomenon of coercion in relation to the concept of affordance. The concept of perceptual affordances can be traced to the psychological research of Gibson (1979), who proposed that a fundamental characteristic of cognition is an agent’s direct perception of opportunities for action in a particular environmental situation. By relying on the notion of affordance, we support a notion of language that, unlike the traditional symbolic view, is grounded in people’s experience in their physical environment, particularly in opportunities for taking actions on objects. This approach creates a framework for interpreting the resolution of coercion as a phenomenon of semantic adjustment that relies on available affordances of objects. In fact, under this view, the possibility itself of implying a covert event in language use is understood as triggered by available affordances of objects, and the task of retrieving this implicit piece of information resides in identifying the specific affordance at play in the surrounding context.

Note that this proposal does not depart from the view that there is stereotypical information associated with lexical entries in terms of *qualia* or other means. Rather, starting with the idea that the default information encoded in words is grounded in the most relevant affordances associated with the word denotation, it examines the interaction between lexically specified information and contextually presented affordance induced by a linguistic discourse. In this paper, we focus on coercion involving artefactual objects, as we are primarily interested in modelling goal-oriented behaviour, and report the results of our first experiments towards the goal of developing a distributional model of coercion interpretation using dependency relationships between target words and their modifiers and adjuncts as the basis for generating on-line transformations that project semantic subspaces in which the interpretations of coercive compositions are expected to emerge as salient word-vectors.

3 Base Methodology: Joint Probabilities of Interpretations

Our objective is to develop a model for predicting interpretations of type coercions inherent in compositions consisting of a transitive verbs that expects an object of a certain type coupled with nouns of a different type in the object position. For the purposes of the present research, we consider an *interpretation* to be a verb that can be inserted into a coercive composition in order to resolve the mismatch between the argument type expected by the coercive verb and the type associated with the object:

2 finish the milk → finish *drinking* the milk

3 cancel the train → cancel *running* the train

As a baseline, we consider a methodology involving a probabilistic analysis of the way that both target verbs and nouns co-occur in dependency relationships with other verbs and verbals. Specifically, given

¹Such cases of coercion to events are also referred to as *logical metonymies* (see Lapata and Lascarides, 2003).

²Recent work in Generative Lexicon claims that the *qualia* value may be updated in context and that the reconstructed event may be assigned contextually (Pustejovsky and Jezek, 2012). For example in the corpus fragment, “So unless the winemakers add tannin by finishing the wine in oak ...”, the context words *winemaker*, *tannin* and *oak* trigger a different interpretation for *wine* (preparing, making) as the object of *finish*.

a target verb v that takes a target noun n as an object, we consider as candidate interpretations of the composition $v(n)$ the intersection of the set of verbs that are observed to take n as an objective argument and the set of verbals (verbs that act as nouns) that are at the head of verbal noun phrases that are objects of v . So for instance, if both the phrases *the boy drinks milk* and *the girl finished drinking* are observed, *drink* is considered a candidate interpretation for *finished the milk*.

Defining this set of k viable interpretations as R , we propose a straightforward probabilistic mechanism for assigning a score to the general appropriateness of a particular candidate interpretation $r_i \in R$:

$$s(r_i) = \frac{f(n(r_i))}{\sum_{j=1}^k f(n(r_j))} \times \frac{f(v(r_i))}{\sum_{j=1}^k f(v(r_j))} \quad (1)$$

Here $f(n(r_i))$ indicates the frequency at which the verb r_i is observed to take the noun n as an objective argument in a corpus, and $f(v(r_i))$ correspondingly indicates the frequency at which r_i is observed at the head of noun phrases that are objects of v . Thus $s(r_i)$ can be interpreted simply as the joint probability of r_i playing the specified compositional roles with the target noun and verb.

To explore the efficacy of this scoring mechanism, we consider a list of candidate verbs and nouns:

VERBS: *finish, begin, enjoy, hear, prefer, cancel*

NOUNS: *coffee, wine, beer, milk, drink, sandwich, cake, dessert, glass, bottle, car, table, door, ambulance, train, book, newspaper, bell, radio, television*

These verbs have been selected for their tendency to coerce objects, and the nouns denote artefactual objects, which is to say, objects that have been made by humans for a reason and should presumably present a range of affordances. In order to assess the proposed metric, we extract relevant dependency relationships from the March 20, 2018 dump of English language Wikipedia.³ A matrix composing every noun with every verb and then listing the top three interpretations in terms of the metric in Equation 1 is reported in Appendix 1.

Here we find a number of outputs that qualitatively appear to be stereotypical of the expectations for some of the more interpretable coercions. Examples include *begin [drinking] coffee*, *enjoy [driving] car*, *cancel [ordering] sandwich*, and so forth. Other cases, such as *finish [starting] newspaper*, do not seem as natural. The verb *hear* in particular appears to pose a challenge for this methodology, and it is worth noting that this verb is in a sense an outlier in that the interpretation would typically involve the action of the coerced object rather than the act of the subject upon that object: *hear* is a perception verb, and so takes an experiencer rather than an agent as a subject. So, for instance, in the phrase *hear the bell [ringing]*, it is what the bell is doing, rather than what the hearer is doing, that is coerced by *hear*. There is also evidence of sense ambiguity, for instance in the case of *table*, which is across the board interpreted as something that a subject would *see*: this is presumably an artefact of a high number of mentions of tables in the sense encountered (“seen”) in a document in our corpus.⁴

4 Extended Methodology: Projections Based on Syntactic Co-Occurrence

The methodology described in the previous section is largely in line with the probabilistic approach proposed by Lapata and Lascarides (2003). As those authors note, this technique does not have a facility for incorporating context into its interpretations: the sentences *the reviewer finished the book* and *the author finished the book* would be interpreted with no consideration of the different actions the respective subjects might take on a book. In the context of affordances, we can say that a book affords different sets of actions to reviewers versus authors.

A straightforward solution to this problem of contextualisation is to introduce a term involving probabilities of observing subjective arguments for candidate verbs. Compounding the score described in Equation 1 with a term for computing the probability of observing a subjective argument for the interpretive verb r_i serves both to lend context to the metric and to limit the set of candidate interpretations R . Here are three toy examples of the effect this step has on output:

³Parsing is performed using the Spacy module for Python.

⁴Word sense ambiguity in semantic type coercion is addressed in depth by Shutova et al. (2013).

4a (<i>brewer, enjoy, beer</i>)	→	<i>produce</i>	4b (<i>patron, enjoy, beer</i>)	→	<i>drink</i>
5a (<i>author, begin, book</i>)	→	<i>write</i>	5b (<i>reviewer, begin, book</i>)	→	<i>write</i>
6a (<i>baker, finish, cake</i>)	→	<i>take</i>	6b (<i>guest, finish, cake</i>)	→	<i>take</i>

The first pair of interpretations is qualitatively satisfying, but the second pair reveals a shortcoming in the methodology: while we may reasonably expect a book to afford reading rather than writing to a reviewer, writing is nonetheless an activity in which reviewers are categorically involved, so the high probability of observing *write(reviewer)* as a subject-predicate composition is imposed on the interpretation. The third pair illustrates the same contextual failure compounded by a tendency to offer overly general interpretations. Furthermore, instances of sentences where the subject offers such straightforward contextualisation are the exception; far more common are, for instance, pronominal subjects with removed antecedents. Straightforward syntactic heuristics are unlikely to offer a consistent way of extrapolating contexts such as subjects from sentences, to the degree that context is available at all.

To address these problems, we propose a methodology that provides a more general framework for using distributional information to generate transformations that can be performed upon a base set of representations for candidate interpretations. We begin by noting that certain grammatical classes such as adjectives and prepositions are suited to convey the affordances associated with particular objects: that a *book* might be *long* or *short*, for instance, or that something might happen *before* or *after* a *beer* suggests an aspectual quality to the actions afforded by these objects.

Following on this observation, we propose a general methodology that involves the construction of tensors based on the probability distributions associated with the co-occurrences of words in certain dependency relationships with input words.⁵ These context-specific tensors can then be used to project a base set of distributional semantic representations into a subspace that captures the salient properties of a particular coercion. This idea is motivated by insight from compositional distributional semantics, where for instance parts of speech such as adjectives can be represented as matrices that transform noun-vectors to a modified representation (Baroni and Zamparelli, 2010), or where sequences of linear algebraic operations over the elements of a sentence map from word-level representations to sentence-level representations (Coecke et al., 2011).

Our conjecture is that these transformations will offer a mechanism for quantitatively encoding the affordances activated by coercions. Given a set of candidate interpretations, we can represent them in terms of their general distributional profile across a corpus (so, for instance, as a set of *base word-vectors* as described by McGregor et al. (2015)). Our objective is to transform these representations of candidate interpretations using the tensor generated by an analysis of dependency relationships, sifting the base space into a contextualised subspace. We hypothesise that geometric characteristics of the context specific subspaces will indicate apt interpretations. The norm of a projected vector, for instance, would be enhanced by the high element-wise overlap with the salient features used to define the transformation tensor, so we would expect good interpretations to drift to the outer fringe of a subspace. In order to illustrate this proposal, we offer a proof-of-concept of how our methodology can work.

A Proof of Concept For an input subject-verb-object tuple (s, v, n) , we consider the intersection A of all adjectives that are observed to co-occur as either a modifier of n or as a modifier of a direct object taken as an argument of the head of s . We construct two probability distributions s_A and n_A based on the L1 normalisations of the frequencies with which the words in A are observed in these relationships with s and n . We compose a vector of joint probabilities $j = \{s_{A_i} \times n_{A_i}\}$ for all $A_i \in A$ and choose the top k adjectives from A based on this metric (for the purposes of the examples here, k is set to 80). If we consider the example of $(\textit{brewer}, \textit{enjoy}, \textit{beer})$, then the six top scoring components that become part of A are *own*, *first*, *strong*, *new*, *local*, and *bottled*, while adjectives such as *traditional*, *seasonal*, *dry*, and *fermented* also appear in the 80 most probable co-occurrences. We extrapolate two new k -dimensional L1 normalised vectors for this set of salient adjectives, s_A^k and n_A^k , and then use the outer product of these vectors to construct a $k \times k$ tensor $T = s_A^k \otimes n_A^k$.

⁵Here *input words* might include a coercive dyad (such as *finish beer*), other words from a phrase or sentence, or indeed topical interpretations of a sentence, in line with the model described by Chersoni et al. (2017).

We consider the vocabulary of target interpretations I to be the intersection of all verbs that are observed to occur as both heads of n and heads of verbal noun phrases that are arguments of v , as in the methodology and examples provided in Section 3, with all verbs observed to take s as a subject. We extrapolate a set of word-vectors W_I corresponding to this vocabulary based on simple co-occurrences statistics for this vocabulary, following the skewed pointwise mutual information weighting scheme described by McGregor et al. (2015), which has been designed for the purpose of projecting conceptually contextualised subspaces of semantic representations. So, for instance, for *(brewer, enjoy, beer)*, the verb *brew* is specified as an argument observed both for *brewer* and at the head of phrases that are arguments of *enjoy*: we therefore include in W_I a vector consisting of features representing the weighting between *brew* and *own, first, new, strong*, and so on. Each vector $w_i \in W_I$ is defined in terms of the same top k elements of A that delineate T . We use T to project each word-vector in W_I into a contextualised subspace $Z = \{Tw_i : w_i \in W_I\}$.

Because the elements of T and W_I are aligned, we expect vectors with co-occurrence features that strongly correlate with the adjectives most salient to the affordances offered by n to s and activated by the coercion of n by v to emerge in Z . We measure this emergence by computing the norm for each transformed vector $z_i \in Z$ and return the vector with the highest norm as our interpretations of (s, v, n) . Results for this methodology applied to the same sentences analysed above play out as follows:

7a <i>(brewer, enjoy, beer)</i> → <i>brew</i>	7b <i>(patron, enjoy, beer)</i> → <i>drink</i>
8a <i>(author, begin, book)</i> → <i>republish</i>	8b <i>(reviewer, begin, book)</i> → <i>bemoan</i>
9a <i>(baker, finish, cake)</i> → <i>bake</i>	9b <i>(guest, finish, cake)</i> → <i>eat</i>

Here we see how this methodology can inject a greater deal of contextuality into its interpretations than the purely probabilistic technique outlined at the beginning of the section. The first and third pairs are categorically plausible interpretations. The second pair is perhaps a bit stranger, with *the reviewer begins [bemoaning] the book* arguably an instance of interpreting one coercion with another, but they are illustrative of the way that a projection of candidate interpretations can expose and exploit the semantic nuance that arises in a communicative context.

5 Conclusion and Continuation

The work presented in this paper is a first attempt at developing a distributional model of coercion interpretation grounded in the theory of perceptual affordances. The idea proposed and provisionally illustrated here is that affordances can be, in a sense, simulated through an analysis of dependency relationships as observed over a large-scale corpus. This analysis has served as the basis for exploring a few examples in which interpretations are conceived of as context-specific projections from a base space of candidate word-vectors. As a first approximation, we have taken adjectives as representative of the properties afforded by particular objects in certain situations, but there is clearly scope for significant expansion of this basic assumption: we might consider other relationships observed with the coerced noun as well as head nouns in subjective roles, words observed in relationships with the coercive predicate (prepositions and intentional adverbs seem intuitively like they might be of interest here), and also words extracted using heuristics outside of dependency parsing.

It is also worth noting that an objective of our modelling approach is to provide a mechanism for representing the overall conceptual context of a sentence. Topic modelling techniques might offer a more holistic framework for the extraction of semantic cues from linguistic data, and there are various options available in this respect. A next step in this project will be the development of a dataset of sentences designed to exhibit various aspects of type coercion, and there is existing work to be considered here, for instance the dataset described by Pustejovsky et al. (2010). The development of a dataset will provide impetus for further refinement of the model itself, as there are clearly a number of directions in which this research can progress.

Acknowledgements

This work has been supported by the CHIST-ERA project ATLANTIS.

References

- Baroni, M. and R. Zamparelli (2010). Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pp. 1183–1193.
- Chersoni, E., A. Lenci, and P. Blache (2017). Logical metonymy in a distributional model of sentence comprehension. In *Sixth Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pp. 168–177.
- Clark, S. (2015). Vector space models of lexical meaning. In S. Lappin and C. Fox (Eds.), *The Handbook of Contemporary Semantic Theory*, pp. 493–522. Wiley-Blackwell.
- Coecke, B., M. Sadrzadeh, and S. Clark (2011). Mathematical foundations for a compositional distributed model of meaning. *Linguistic Analysis* 36(1-4), 345–384.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Greenberg, C., A. Sayeed, and V. Demberg (2015). Improving unsupervised vector-space thematic fit evaluation via role-filler prototype clustering. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 21–31.
- Lapata, M. and A. Lascarides (2003). A probabilistic account of logical metonymy. *Computational Linguistics* 29(2), 261–315.
- McGregor, S., K. Agres, M. Purver, and G. Wiggins (2015). From distributional semantics to conceptual spaces: A novel computational method for concept creation. *Journal of Artificial General Intelligence* 6(1), 55–89.
- McGregor, S. and K. Lim (2018). Affordances in grounded language learning. In *Proceedings of the Eight Workshop on Cognitive Aspects of Computational Language Learning and Processing*.
- Moens, M. and M. Steedman (1988). Temporal ontology and temporal reference. *Computational linguistics* 14(2), 15–28.
- Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics* 17(4), 409–441.
- Pustejovsky, J. and E. Jezek (2012). *Introducing Qualia Structure*. Manuscript, Brandeis University and University of Pavia, pp. 1–45. Available at <https://http://gl-tutorials.org/wp-content/uploads/2015/12/GL-QualiaStructure.pdf>.
- Pustejovsky, J., A. Rumshisky, A. Plotnick, E. Jezek, O. Batiukova, and V. Quochi (2010). SemEval-2010 Task 7: Argument Selection and Coercion. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pp. 27–32.
- Raczaszek-Leonardi, J., I. Nomikou, K. J. Rohlfing, and T. W. Deacon (2018). Language development from an ecological perspective: Ecologically valid ways to abstract symbols. *Ecological Psychology* 30(1), 39–73.
- Shutova, E., J. Kaplan, S. Teufel, and A. Korhonen (2013). A computational model of logical metonymy. *ACM Transactions on Speech and Language Processing* 10(3), 11:1–11:28.
- Zarcone, A., A. Lenci, S. Padó, and J. Utt (2013). Fitting, not clashing! a distributional semantic model of logical metonymy. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)*, pp. 404–410.

Appendix 1

	finish	begin	enjoy	hear	prefer	cancel
coffee	serve take produce	drink grow serve	drink get serve	get include give	grow serve drink	produce include order
wine	produce take drink	produce drink sell	produce drink create	produce give include	produce drink give	produce give create
beer	produce serve drink	produce drink sell	drink produce create	produce give sell	produce drink sell	produce give order
milk	take produce give	take produce drink	drink take produce	take give include	take give produce	produce take give
drink	take serve finish	take get serve	take give get	give take get	take give buy	take give order
sandwich	eat serve take	eat sell offer	eat create get	include get give	eat sell serve	order include offer
cake	take eat win	take eat produce	eat create take	give take include	eat take give	give take produce
dessert	win serve eat	create produce prepare	create eat play	include call play	create eat call	create produce include
glass	take produce read	wear take produce	wear read drink	include give take	wear take give	produce take give
bottle	take finish produce	take produce drink	drink take give	give take find	take give keep	produce give take
car	build take drive	take build produce	drive take see	take drive include	take build drive	build produce take
table	see finish lead	see turn take	see play create	see include turn	see leave keep	see turn leave
door	open go close	open close break	open go work	open go close	open keep leave	open close leave
ambulance	call drive take	call operate drive	call drive see	call include drive	call drive see	call meet drive
train	take run work	take run operate	work take run	take run see	take work run	meet take run
book	write read publish	write publish read	write read publish	write read call	write call publish	publish release produce
newspaper	start read write	publish tell sell	read tell write	tell read start	publish read leave	publish leave start
bell	play take tie	play take ring	play hear give	ring play say	play take call	cast hear hold
radio	build play take	take turn play	play do work	talk play include	leave play take	leave include build
television	record watch write	watch broadcast take	watch do play	watch do record	watch leave do	leave watch produce

Table 1: The top three interpretations for all compositions of 6 coercive verbs with 19 artefactual objects, based on the joint probability of the interpretation occurring with each input word. The outputs *have*, *make*, and *use* are almost ubiquitously returned, and have been suppressed here.