A Type-Driven Vector Semantics for Ellipsis with Anaphora Using Lambek Calculus with Limited Contraction

Gijs Wijnholds¹ · Mehrnoosh Sadrzadeh¹

Published online: 20 May 2019
© The Author(s) 2019

Abstract

We develop a vector space semantics for verb phrase ellipsis with anaphora using type-driven compositional distributional semantics based on the Lambek calculus with limited contraction (LCC) of Jäger (Anaphora and type logical grammar, Springer, Berlin, 2006). Distributional semantics has a lot to say about the statistical collocation based meanings of content words, but provides little guidance on how to treat function words. Formal semantics on the other hand, has powerful mechanisms for dealing with relative pronouns, coordinators, and the like. Type-driven compositional distributional semantics brings these two models together. We review previous compositional distributional models of relative pronouns, coordination and a restricted account of ellipsis in the DisCoCat framework of Coecke et al. (Mathematical foundations for a compositional distributional model of meaning, 2010. [arXiv:1003.4394](#), Ann Pure Appl Log 164(11):1079–1100, 2013). We show how DisCoCat cannot deal with general forms of ellipsis, which rely on copying of information, and develop a novel way of connecting typelogical grammar to distributional semantics by assigning vector interpretable lambda terms to derivations of LCC in the style of Muskens and Sadrzadeh (in: Amblard, de Groote, Pogodalla, Retoré (eds) Logical aspects of computational linguistics, Springer, Berlin, 2016). What follows is an account of (verb phrase) ellipsis in which word meanings can be copied: the meaning of a sentence is now a program with non-linear access to individual word embeddings. We present the theoretical setting, work out examples, and demonstrate our results with a state of the art distributional model on an extended verb disambiguation dataset.

Keywords Compositional distributional semantics · Typelogical grammar · Ellipsis · Non-linearity · Lambda calculus

✉ Gijs Wijnholds
g.j.wijnholds@qmul.ac.uk
Mehrnoosh Sadrzadeh
mehrnoosh.sadrzadeh@qmul.ac.uk

¹ School of Electronic Engineering and Computer Science, Queen Mary University of London, London, UK

1 Introduction

Distributional semantics is a field of research within computational linguistics that provides an easily implementable algorithm with an empirically verifiable output for representing word meanings and degrees of semantic similarity thereof. This semantics is rooted in the distributional hypothesis, often referred to via the quote “you shall know a word by the company it keeps”, made popular by Firth (1957). More precisely, according to the distributional hypothesis words that occur in similar contexts have similar meaning. This idea has been made concrete by gathering the co-occurrence statistics of context and target words in corpora of text and using that as a basis for developing vector representations for word meanings. A notion of similarity based on the cosine of the angle between vectors allows one to compare degrees of word similarity in the vector space models where these word vectors embed. Such models have been shown to perform well in a variety of natural language processing (NLP) tasks, such as semantic priming (Lund and Burgess 1996) and word sense disambiguation (Schütze 1998). The underlying philosophy has gained attention in cognitive science as well (Lenci 2008).

Although this notion of similarity is intuitive and works well at the word level, it is less productive to consider phrases and full sentences to be similar whenever they occur in a similar context. Firstly, we know that language is compositional, since the number of potential sentences humans can produce is larger than the amount a single human ever produces. Secondly, data sparsity issues arise when treating sentences as individual expressions and computing direct co-occurrence statistics for them. So the challenge of producing vector representations for phrases and sentences rests on the shoulders of compositional distributional semantics. Several studies have tried to learn not just vectors for words, but embeddings for several constituents (Baroni and Zamparelli 2010; Grefenstette et al. 2013), or have experimented with simple commutative compositional operations such as addition and multiplication (Mitchell and Lapata 2008). A structured attempt at providing a general mathematically sound model of compositional distributional semantics has been presented by Coecke et al. (2010); these models start from the observation that vector spaces share the same structure as Lambek’s most recent grammar formalism, pregroup grammar (Lambek 1997), and interpret its derivations in terms of vector spaces and linear maps. What follows is an architecture that is familiar from logical formal semantics in Montague style (Montague 1970), where the judgments of a grammar translate to a consistent semantic operation (read linear map) that acts on the individual word vectors to produce some vector in the sentence space. A number of subsequent attempts has shown that a similar interpretation is possible for other typological grammars, such as Lambek’s original syntactic calculus (Coecke et al. 2013), Lambek–Grishin grammars (Wijnholds 2014), and Combinatorial Categorical Grammar (CCG) (Maillard et al. 2014).

One major issue for distributional semantics and especially compositional approaches therein is to find a suitable representation for function words. Without the power of formal semantics to assign constant meanings or to allow set-theoretic operations, distributional semantics does not have much to say about the meaning of logical words such as ‘and’, ‘despite’ and pronouns like ‘his’, ‘which’, ‘that’, let alone quantificational constituents (‘all’, ‘some’, ‘more than half’). All of these words have

in common that they intuitively do not bear a contextual meaning: a function word may co-occur with any content word and so its distribution does not reveal much about its meaning, unless perhaps the notion of meaning is taken to be conversational.¹ To overcome this issue, Sadrzadeh et al. (2013) relies on Frobenius algebras to formalise the notion of *combining* and *dispatching* of information. This approach has seen applications to relative pronouns (Sadrzadeh et al. 2013, 2014), coordination (Kartsaklis 2016), and to a lesser extent to some limited forms of ellipsis (Kartsaklis et al. 2016). In each of these, the Frobenius algebras allow one to use element wise multiplication of arbitrary tensors, corresponding to the usual intersective interpretation one finds in formal semantics (Dalrymple et al. 1991). A treatment of quantification was also given using the bialgebraic nature of vector spaces over powersets of elements (Hedges and Sadrzadeh 2016; Sadrzadeh 2016). An explanation of the derivational processes resulting in these compositional meanings requires more elaborate grammatical mechanisms: Wijnholds (2014) repeats the exercise to give a compositional distributional model for a symmetric extension of the Lambek calculus. A derivational account of pronoun relativisation in English and Dutch is given by means of a Lambek grammar with controlled forms of movement and permutation in Moortgat and Wijnholds (2017), whereas Wijnholds (2019) provides a derivational account of quantification based on the bialgebra account of Hedges and Sadrzadeh 2016; Sadrzadeh 2016.

In this paper, we contribute to the typological style of compositional distributional semantics by giving an account for verb phrase ellipsis with anaphora in a revision of the framework described above. The case of ellipsis traditionally has been approached both as a syntactic problem within categorial grammars (Kubota and Levine 2017; Jäger 2006, 1998; Morrill and Merenciano Saladrigas 1996; Hendriks 1995) as well as as a semantic problem by directly appealing to their lambda calculi term logics (Dalrymple et al. 1991; Kempson et al. 2015). The research within categorial grammar either suggests that elliptic phenomena should be treated using a specific controlled form of copying of information at the antecedent and moving it to the site of ellipsis, e.g. in Jäger (1998) and Morrill and Merenciano Saladrigas (1996), or by maintaining a non-directional functional type (meaning that it is not sensitive to where its argument occurs, before or after it), which is backward/forward looking, e.g. in Kubota and Levine (2017) and Jäger (2006). The first proposal can also be implemented using different modal Lambek Calculi, e.g. that developed in Moortgat (1997) and the second one using Displacement Calculus (Morrill and Valentin 2010); Abstract Categorial Grammars of Muskens (2003) and de Groote (2001), which allow for a separation of syntax and semantics within a categorial grammar and allow for freedom of copying and movement at the semantic side, can also be employed. We will not go too much into philosophical discussion in this paper and base our work on an extension of the Lambek calculus with a limited form of contraction (shorthand to **LCC**) via a non-directional functional type, introduced by (Jäger 2006). In a previous paper (Wijnholds and Sadrzadeh 2018) we showed how one can treat ellipsis using a controlled form of copying and movement via contraction and a modality. What was novel in previous work was that we discovered and showed how the use of Frobenius copying/dispatching of information does not work for resolving ellipsis,

¹ The work of Kruszewski et al. (2016) gives a distributional semantic account of conversational negation.

as it cannot distinguish between sloppy and strict readings. Similar to previous work (Wijnholds and Sadrzadeh 2018), we argue for a simple revision of the DisCoCat framework (Coecke et al. 2013) in order to allow us to incorporate a proper notion of reuse of resources: instead of directly interpreting derivations as linear maps, where it becomes impossible to have a map that copies word embeddings (Jacobs 2011; Abramsky 2009), we decompose the interpretation of grammar derivations into a two-step process, relying on a non-linear simply typed lambda calculus, in the style of (Muskens and Sadrzadeh 2016, 2019). In doing so, we obtain a model that allows for the reuse of embeddings, while staying in the realm of vector spaces and linear maps. The novel part of the current paper, apart from its use of a backward looking bidirectional operation in Lambek Calculus, rather than copying and moving syntactic information around, is that we test our hypothesis on a well known verb disambiguation task (Mitchell and Lapata 2008; Grefenstette and Sadrzadeh 2011a; Kartsaklis and Sadrzadeh 2013) using vectors and matrices obtained from large scale data.

The paper is structured as follows: Sect. 2 discusses the problem of ellipsis and anaphora and argues for non-linearity in the syntactic process, Sect. 3 gives the general architecture of our system. We proceed in Sect. 4 with our main analysis and carry out a simple experiment in Sect. 5 to show how our model may be empirically validated. We conclude with a discussion and future work in Sect. 6.

2 Ellipsis and Non-linearity

Ellipsis can be defined as a linguistic phenomenon in which the full content of a sentence differs from its representation. In other words, in a case of ellipsis a phrase is *missing* some part needed to recover its meaning. There are numerous types of ellipsis with a varying degree of complexity, but we will stick with verb phrase ellipsis, in which very often an ellipsis *marker* is present to mark what part of the sentence is missing and where it has to be placed.

An example of verb phrase ellipsis is in Eq. 1, where the elided verb phrase is marked by the auxiliary verb. Ideally, sentence 1(a) is in a bidirectional entailment relation with 1(b), i.e. (a) entails (b) and (b) entails (a).

$$\begin{array}{ll} a \text{ "Alice drinks and Bill does too"} & \\ b \text{ "Alice drinks and Bill drinks"} & (1) \end{array}$$

More complicated examples of ellipsis introduce an ambiguity; the example in Eq. 2, has a sloppy (b) and a strict (c) interpretation for (a).

$$\begin{array}{ll} a \text{ "Gary loves his code and Bill does too"} & \text{(ambiguous)} \\ b \text{ "Gary loves Gary's code and Bill loves Bill's code"} & \text{(sloppy)} \\ c \text{ "Gary loves Gary's code and Bill loves Gary's code"} & \text{(strict)} \end{array} \quad (2)$$

In a formal semantics account, the first example could be analysed with the auxiliary verb as an identity function on the main verb of the sentence and an intersective meaning for the coordinator. Somehow the parts need to be appropriately combined to produce the reading (b):

doestoo : $\lambda x.x$
 and : $\lambda x.\lambda y.(x \wedge y)$ should give **drinks(alice) $\wedge$ drinks(bill)**

The second example would assume the same meaning for the coordinator and auxiliary but now the possessive pronoun “his” gets a more complicated term: $\lambda x.\lambda y.\text{owns}(x, y)$. The analysis then somehow should derive two readings:

loves(gary, x) $\wedge$ owns(gary, x) $\wedge$ loves(bill, x)
loves(gary, x) $\wedge$ owns(gary, x) $\wedge$ loves(bill, y) $\wedge$ owns(bill, y)

Indeed, these are meanings that would be produced by the approach of Dalrymple et al. (1991).

There are three issues with these analyses that are not solved in current distributional semantic frameworks: first, it is unclear what the composition operator is that maps from the meanings of the words to the meaning of the phrases. Second, it is unclear how the lexical constants (mainly the intersection operation expressed by the conjunction $\wedge$) are to be interpreted as a linear map. Thirdly, these examples contain a non-linearity; resources may be used more than once (the main verb is used twice in the first example, the noun phrases in the second example). We will outline a model that deals with all three.

The challenge of composition will be treated by using a compositional distributional semantic model in the lines of Coecke et al. (2013). For the interpretation on the lexical level, we will make use of Frobenius algebras to specify the lexical meaning of the coordinator and relative pronoun, following Kartsaklis (2016) and Sadrzadeh et al. (2013), respectively. Moreover, we will assign a similar meaning to the possessive pronoun.

What remains is to decide how to deal with non-linearity. With non-linearity we mean the possible duplication of resources, not the use of non-linear maps. On the side of vector semantics, though, it can be tempting to rely on the use of Frobenius algebras and use their dispatching operation to deal with copying, and indeed this operation has been referred to as copying in the literature, e.g. see Coecke and Paquette (2008), Coecke et al. (2008), Kartsaklis (2016) and Sadrzadeh et al. (2013). Going a bit deeper, however, reveals that this operation places a vector into the diagonal of a matrix, that is, for a finite dimensional vector space W spanned by basis $\{\vec{b}_i\}_i$, we have

$$\Delta: W \rightarrow W \otimes W \quad \text{given by} \quad \Delta\left(\sum_i C_i \vec{b}_i\right) = \sum_{ii} C_i \vec{b}_i \otimes \vec{b}_i$$

As an example consider a two dimensional space W . A vector in this space will be copied via Δ into a matrix in $W \otimes W$, whose diagonals are a and b and whose non-diagonals are 0:

$$\begin{pmatrix} a \\ b \end{pmatrix} \xrightarrow{\Delta} \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$$

This is computed from the definition of Δ on the basis $\Delta(\vec{b}_i) = \vec{b}_i \otimes \vec{b}_i$ and the fact that Δ is linear. Indeed it seems that Δ only “copies” the basis into their tensors and it has been shown that any other form of copying in this context, e.g. a Cartesian one, is not allowed, see Jacobs (2011) and Abramsky (2009) for proofs. For a concrete linguistic demonstration of this fact, consider the anaphoric sentence “Alice loves herself”, with the following noun vector and verb tensor:

$$\mathbf{alice} = \sum_i a_i \mathbf{v}_i \quad \mathbf{loves} = \sum_{jkl} c_{jkl} (\mathbf{v}_j \otimes \mathbf{s}_k \otimes \mathbf{v}_l)$$

We want to obtain the interpretation

$$\mathbf{alice}_i \mathbf{loves}_{ijk} \mathbf{alice}_k = \sum_{ijk} a_i c_{ijk} a_k \mathbf{s}_j$$

However, using Frobenius copying would not give the desired result but rather something else:

$$\mathbf{alice}_i \mathbf{loves}_{iji} = \sum_{ij} a_i a_i c_{iji} \mathbf{s}_j$$

Note that in the above, we use a simplification of Einstein’s index notation for tensors. In Einstein notation, a tensor has indices on the top and bottom, specifying which index refers to a row or a column. For instance, a matrix is denoted by ${}^i M_j$, when i enumerates the row elements and j the column elements. We, however, only work with finite dimensional vector spaces where a space is isomorphic to its dual space. In such cases, the Einstein notation simplifies and one can write both of the subscripts under (or above).

The use of Frobenius operations in itself may not immediately appear to be a problem, but in previous work (Wijnholds and Sadrzadeh 2018) we show how a categorical model in the framework of Coecke et al. (2013), using Frobenius operations as a ‘copying’ operation, and treating the auxiliary phrase “does too” as an identity map, is not able to distinguish between the sloppy versus strict interpretations of a sentence with an elliptical phrase such as “Alice loves her code, so does Mary”. The vector interpretations of both of the cases (1) “Alice loves Alice’s code and Mary loves Alice’s code” and (2) “Alice loves Alice’s code and Mary loves Mary’s code” become the following expression:

$$\begin{aligned} & \mathbf{Alice} \otimes \mathbf{loves} \otimes \mathbf{her} \otimes \mathbf{code} \otimes \mathbf{and} \otimes \mathbf{Mary} \otimes \mathbf{does\ too} \\ & \mapsto \Delta_N(\mathbf{Alice} \odot \mathbf{Mary} \odot \mathbf{code})_{ik} \mathbf{loves}_{ijk} \end{aligned}$$

Even though one may attempt to fix this problematic behaviour by complicating the meaning of an auxiliary verb, it shows that under reasonable assumptions a direct translation of proofs into a Frobenius tensor based model is not desirable for approaches that require this non-linear behaviour. In order to still have Cartesian copying behaviour in

a tensor based model, we decompose the DisCoCat model of Coecke et al. (2013) into a two-step architecture: we first define an extension of the Lambek calculus, which allows for limited contraction, developed in Jäger (2006). In this setting, grammatical derivations as well as lexical entries are interpreted in a non-linear simply typed λ -calculus. The second stage of interpretation homomorphically maps these abstract meaning terms to terms in a lambda calculus of vectors, tensors, and linear maps, developed in Muskens and Sadrzadeh (2016). The final effect is that we allow the Cartesian behavior of copying elements before concretisation in a vector semantics: the meaning of a sentence now is a program that has non-linear access to word embeddings.

3 Typological Distributional Semantics

In a very general setting, compositionality can be defined as a homomorphic image (or functorial passage) from a syntactic algebra (or category) to a semantic algebra (or category). The only condition, then, is that the semantic algebra be weaker than the syntactic algebra: each syntactic operation needs to be interpretable by a semantic operation. To give a formal semantic account one would map the proof terms of a categorial grammar, or rewritings of a generative grammar, to the semantic operations of abstraction and application of some lambda term calculus. In a distributional model such as the one of Coecke et al. (2010, 2013), derivations of a Lambek grammar are interpreted by linear maps on finite dimensional vector spaces. For our presentation it will suffice to say that the Lambek calculus can be considered to be a *monoidal biclosed category*, which makes the mapping to the compact closure of vector spaces straightforward. However, we want to employ the copying power of non-linear lambda calculus, and so we will move from the direct interpretation below

to a two-step interpretation process:

What we end up with is in fact a more intricate target than in the direct case: target expressions are now lambda terms with a tensorial interpretation, i.e. a program with access to word embeddings. The next subsections outline the details: we consider the syntax of the Lambek calculus with limited contraction, the semantics of non-linear lambda calculus, and the interpretation of lambda terms in a lambda calculus with tensors.

3.1 Derivational Semantics: Formulas, Proofs and Terms

We start by introducing the Lambek Calculus with Limited Contraction **LLC**, a conservative extension of the Lambek calculus **L**, as defined in Jäger's monograph (Jäger 2006).

LLC was in first instance defined to deal with anaphoric binding, but has many more applications including verb phrase ellipsis and ellipsis with anaphora. The system extends the Lambek calculus with a single binary connective $|$ that behaves like an implication for anaphora: a formula $A|B$ says that a formula A can be bound to produce formula B , while retaining the formula A . This non-linear behaviour allows the kind of resource multiplication in syntax that one expects when dealing with anaphoric binding and ellipsis.

More formally, formulas or types of **LLC** are built given a set of basic types T and using the following definition:

Definition 1 Formulas of **LLC** are given by the following grammar, where T is a finite set of basic formulas:

$$A, B := T \mid A \bullet B \mid A \backslash B \mid B / A \mid A|B$$

Intuitively, the Lambek connectives $\bullet$, $\backslash$, $/$ represent a 'logic of concatenation': $\bullet$ represents the concatenation of resources, where $\backslash$, $/$ represent directional decatenation, behaving as the residual implications with respect to the multiplication $\bullet$. The extra connective $|$ is a separate implication that behaves non-linearly: its deduction rules allow a mix of permutations and contractions, which effectively treat anaphora and VP ellipsis markers as phrases that look leftward to find a proper binding antecedent. Our convention is that we read $A|B$ as a function with input of type A and an output of type B .

The rules of **LLC** are given in a natural deduction style in Fig. 1. The *Lex* rule is an axiom of the logic: it allows us to relate the judgements of the logic to the words of the lexicon. For instance, in the example proof tree provided in Fig. 2, the judgement *alice* : *np* is related to the word *Alice*, the judgement *bob* : *np* to the word *Bob*, and the judgement *and* : $(s \backslash s)/s$ to the word *and*. Then, as it is usual in natural deduction, every connective has an introduction rule, marked with *I* and an elimination rule, marked with *E*. In the introduction rules for $/$ and $\backslash$, the variable x stands for an axiom, in the introduction rule for $\bullet$ and eliminations rules for $\bullet$, $/$ and $\backslash$, we have proofs for the premise types A , B , $A \bullet B$, A/B and $A \backslash B$, i.e. general terms N and M .

Informally speaking, the introduction rule for $\bullet$, takes two terms M and N , one of which (M) proves a formula A and another of which (N) proves the formula B , and it pairs the terms with the tensor product of the formulae, that is, it tells us that the $\langle M, N \rangle$ proves $A \bullet B$. The elimination rule for $\bullet$ takes a pair of terms, denoted by M and tells us that the first projection of M , i.e. $\pi_1(M)$, i.e. the first element of the pair, proves A and its second projection/element proves B . The introduction rule for $\backslash$ takes the index of the rule where formula A was proved using a term x , a proof tree which used this rule and possibly together with some other rules proved the formula

Fig. 1 Labelled natural deduction for **LLC**

$$\begin{array}{c}
 \frac{w}{x : \sigma(w)} \text{Lex} \\
 \\
 \frac{M : A \quad N : B}{\langle M, N \rangle : A \bullet B} I \bullet \qquad \frac{M : A \bullet B}{\pi_1(M) : A \quad \pi_2(M) : B} E \bullet \\
 \\
 \frac{\overline{x : A} \quad \begin{array}{c} i \\ \vdots \end{array}}{M : B} \quad \frac{M : B}{\lambda x. M : A \setminus B} I \setminus, i \qquad \frac{N : A \quad M : A \setminus B}{M N : B} E \setminus \\
 \\
 \frac{\overline{x : A} \quad \begin{array}{c} i \\ \vdots \end{array}}{M : B} \quad \frac{M : B}{\lambda x. M : B / A} I /, i \qquad \frac{M : B / A \quad N : A}{M N : B} E / \\
 \\
 [N : A]_i \quad \dots \quad \frac{M : A \setminus B}{M N : B} E \setminus, i
 \end{array}$$

B using the term M , then derives the formula $A \setminus B$ using the lambda term $\lambda x. M$. The lambda terms are explained later on, but for now, think of this term as a function M with the variable x . The elimination rule for $\setminus$ is doing the opposite of what we just explained and which is what the introduction rule did. It takes a term x for formula A , a term y for formula $A \setminus B$, then tells us that we can apply y to x to get something of type B . The rules for $/$ are similar to these but with different ordering, which is easily checkable from their proof rules in Fig. 1.

That brings us to the main rules that differentiate **LLC** from **L** (the Lambek Calculus): the rules for $|$. Here, the elimination rule tells us that if somewhere in the proof we had proved A from N , and denoted the result by an index i , and then later we encounter a term M for $A|B$, and that i happened before $M : A|B$, then we are allowed to eliminate $|$ and get B by applying the term M to the term N . This rule is very similar to either of the $\setminus$ and $/$ rules, in that it says you can eliminate the connective by applying its term to the term of one of its compartments, i.e. its input. The exception for the $|$ elimination rule is that it allows for that input, i.e. $[N : A]_i$ to happen not directly as the antecedent of the elimination rule, but as one of the other rules in the proof, somewhere before the current elimination rule. We can see how this rule is applicable in the proof tree of Fig. 2: we see a proof for $\text{drinks} : np \setminus s$, in this occasion indexed with a label i , then quite later on in the proof (actually at the end of it), we encounter the term $\text{dt drinks} : np \setminus s$, now the $E|, i$ rule allows us to apply the latter to the former, all the way back, to obtain $\text{dt} : (np \setminus n)(np \setminus s)$. The $|$ connective also has an introduction rule, a proper formulation of this rule, however, is more delicate. Since our anaphoric expressions are already typed in the lexicon, we do not need this rule in our paper and refer the reader for different formulations and explanations of it to Jäger's book (Jäger 2006, pp.123–124).

The interpretation of proofs is established by a non-linear simply typed lambda term calculus, which labels the natural deduction rules of the calculus:

Definition 2 Given a countably infinite set of variables $V = \{x, y, z \dots\}$, terms of λ are as in the below grammar:

$$M, N := V \mid \lambda x.M \mid M N \mid \langle M, N \rangle \mid \pi_1(M) \mid \pi_2(M)$$

Terms obey the standard α -, η - and β -conversion rules:

Definition 3 For terms of λ we define three conversion relations:

1. α -conversion: for any term M we have

$$M =_\alpha M[x \mapsto y]$$

provided that y is a fresh variable, i.e. it does not occur in M .

2. η -conversion: for terms M we have

$$\begin{aligned} \lambda x.M x &=_\eta M \quad (x \text{ does not occur in } M) \\ \langle \pi_1(M), \pi_2(M) \rangle &=_\eta M \end{aligned}$$

3. β -conversion: for terms M we define

$$\begin{aligned} (\lambda x.M) N &\rightarrow_\beta M[x \mapsto N] \\ \pi_1(\langle M, N \rangle) &\rightarrow_\beta M \\ \pi_2(\langle M, N \rangle) &\rightarrow_\beta N \end{aligned}$$

We moreover write $M \rightarrow_\beta^* N$ whenever M converts to N in multiple steps.

The full labelled natural deduction is given in Fig. 1. Proofs and terms give the basis of the derivational semantics; given a lexical map relation $\sigma \subseteq \Sigma \times F$ for Σ a dictionary of words, we say that a sequence of words $w_1, \dots, w_n$ derives the formula A whenever it is possible to derive a term $M : A$ with free variables x_i of type $\sigma(w_i)$. For the x_i , one can substitute constants c_i of type $\sigma(w_i)$ representing the meaning of the actual words $w_1, \dots, w_n$. The *abstract meaning* of the sequence is thus given by the lambda term t . An example of such a proof is given for the elliptical phrase “Alice drinks and Bob does-too” in Fig. 2. More involved examples are given in Figs. 3 and 4; they will be discussed in Sect. 4.

3.2 Lexical Semantics: Lambdas, Tensors and Substitution

We complete the vector semantics by adding the second step in the interpretation process, which is the insertion of lexical entries for the assumptions occurring in a proof. In this step, we face the issue that interpretation directly into a vector space is not an option given that there is no copying map that is linear, while at the same time lambda terms don’t seemingly reflect vectors. We solve the issue by showing, following Muskens and Sadrzadeh (2016, 2019), that vectors can be emulated using a lambda calculus.

$$\frac{\frac{\frac{\text{Alice}}{\text{alice} : np} \text{Lex} \quad \frac{\text{drinks}}{[\text{drinks} : np \backslash s]_i} \text{Lex}}{\text{drinks alice} : s} E \backslash \quad \frac{\frac{\text{and}}{\text{and} : (s \backslash s) / s} \text{Lex} \quad \frac{\frac{\text{Bob}}{\text{bob} : np} \text{Lex} \quad \frac{\frac{\text{does too}}{\text{dt} : (np \backslash s) | (np \backslash s)} \text{Lex}}{\text{dt drinks} : np \backslash s} E \backslash, i}{(\text{dt drinks}) \text{ bob} : s} E \backslash}{\text{and } ((\text{dt drinks}) \text{ bob}) : s \backslash s} E /}{(\text{and } ((\text{dt drinks}) \text{ bob}))(\text{drinks alice}) : s} E \backslash$$

Fig. 2 Short hand derivation for “Alice drinks and Bob does-too”

3.2.1 Lambdas and Tensors

The idea of modelling tensors with lambda calculus is to represent vectors as functions from natural numbers to the values in the underlying field. This representation treats vectors as lists of elements of a field, for instance the field of reals $\mathbb{R}$. What the function is doing is enumerating the elements of this list. So for instance, consider the following vector

$$\vec{v} = [a, b, c, \dots] \quad \text{for } a, b, c, \dots \in \mathbb{R}$$

The representation of $\vec{v}$ using a function f becomes as follows

$$f(1) = a, f(2) = b, f(3) = c, \dots \text{ and so on}$$

For natural language applications, it is convenient to work with a fixed set of indices rather than directly working with natural numbers as the starting point. These indices will be the “context words” of the vector space model of word meaning. For demonstration purposes, suppose these context words are the following set of words

$$C = \{\text{human, painting, army, weapon, marathon}\}$$

Then a “target word”, i.e. the word whose meaning we are representing using these context words, will have values from $\mathbb{R}$ in the entries of a vector spaces spanned by the above context set. For instance, consider three target words “warrior”, “sword”, and “athlete”. Their vector representations are as follows:

$$\begin{aligned}
\vec{\text{warrior}} &= [4, 1, 2, 9, 1] \\
\vec{\text{sword}} &= [2, 3, 9, 2, 0] \\
\vec{\text{athlete}} &= [6, 2, 0, 1, 9]
\end{aligned}$$

In a functional notation, our index set is the set of the context words, e.g. C , as given above, and for each target word t , our function returns its value on each of the context words. So for instance, for a function f , the vector representation of “warrior” becomes as follows

$$f(\text{human}) = 4, f(\text{painting}) = 1, f(\text{army}) = 2, f(\text{weapon}) = 9, f(\text{marathon}) = 1$$

Type-theoretically, instead of working with a set of words as the domain of the representation function f , we enumerate the set of context words and use their indices as inputs to f . So we denote our set C above by indices $i_1, i_2, \dots, i_5$, which changes the function representation to the following

$$f(i_1) = 4, f(i_2) = 1, f(i_3) = 2, f(i_4) = 9, f(i_5) = 1$$

That is, for any dimensionality n , we assume a basic type I_n , representing a finite index set (in concrete models the number of index types will be finite) of context words. The underlying field, in the case of natural language applications remains the set of real numbers $\mathbb{R}$; we denote it by the type R . For more information about $\mathbb{R}$ as a type, see Muskens and Sadrzadeh (2016). As explained above, the type of a vector in $\mathbb{R}^n$ becomes $V^n = I_n \rightarrow R$. Similarly, the type of an $n \times m$ matrix, which is vector in a space whose basis are pairs of words, is $M^{n \times m} = I_n \rightarrow I_m \rightarrow R$. In general, we may represent an arbitrary tensor with dimensions $n, m, \dots, p$ by $T^{n \times m \times \dots \times p} = I_n \rightarrow I_m \rightarrow \dots \rightarrow I_p \rightarrow R$. We abbreviate cubes $T^{n \times m \times p}$ to C and hypercubes $T^{n \times m \times p \times q}$ to H . We will leave out the superscripts denoting dimensionality when they are either irrelevant or understood from the context.

By reference to index notation for linear algebra, we write v i as v_i whenever it is understood that i is of type I . We moreover assume constants for the basic operations of a vector space: $0 : R, 1 : R, + : R \rightarrow R \rightarrow R, \cdot : R \rightarrow R \rightarrow R$ with their standard interpretation. Some standard operations can now be expressed using lambda terms: We can also express many other operations in the same, e.g.

Name	Symbol	Lambda term
Matrix transposition	T	$\lambda m i j . m_{ji} : M \rightarrow M$
Matrix-Vector multiplication	$\times_1$	$\lambda m v i . \sum_j m_{ij} \cdot v_j : M \rightarrow V \rightarrow V$
Cube-Vector multiplication	$\times_2$	$\lambda c v i j . \sum_k c_{ijk} \cdot v_k : C \rightarrow V \rightarrow M$
Hypercube-Matrix multiplication	$\times_3$	$\lambda c m i j . \sum_l c_{ijkl} \cdot m_{kl} : H \rightarrow M \rightarrow M$
Vector Element wise multiplication	$\odot$	$\lambda u v i . u_i \cdot v_i : V \rightarrow V \rightarrow V$
Vector addition	$+$	$\lambda u v i . u_i + v_i : V \rightarrow V \rightarrow V$

backwards matrix multiplication by composing matrix transposition with standard multiplication: $\times^T := \lambda m v i . \sum_j m_{ji} \cdot v_j : M \rightarrow V \rightarrow V$. In the same way, it is

routine to define a cube-matrix multiplication and a hypercube-cube and hypercube-vector multiplication. These operations do not occur in the current paper. Similarly, one can define addition and element wise multiplication operations between matrices, cubes, and hypercubes. In what follows, we abuse the notation and denote the latter two with the same symbols, that is with $+$ and $\odot$ regardless of the type of object they are adding or multiplying.

All of these operations, except for addition, are instances of the multilinear algebraic operation of tensor contraction applicable to any two tensors of arbitrary rank as long

Table 1 Translation that sends abstract terms to a tensor-based model using matrix and cube multiplication as the main operations; here an in the two other proceeding tables the atomic types are np and s

w	$\sigma(w)$	$\mathcal{H}(w)$	$\mathcal{T}(w)$
cn	n	cn	V
adj	n/n	$\lambda v.(\mathbf{adj} \times_1 v)$	VV
adv	$(np \setminus s) \setminus (np \setminus s)$	$\lambda m.(\mathbf{adv} \times_3 m)$	HM
itv	$np \setminus s$	$\lambda v.(\mathbf{itv} \times_1 v)$	VV
tv	$(np \setminus s)/np$	$\lambda uv.(\mathbf{tv} \times_2 u) \times_1 v$	VVV
coord	$(s \setminus s)/s$	$\lambda P.\lambda Q.P \odot Q$	VVV

as they share at least one index. The tensor contraction between them is formed by applying the following formula:

$$\sum_{i_1, \dots, i_{n+k}} A_{i_1 i_2 \dots i_n} B_{i_n i_{n+1} \dots i_{n+k}} \in \underbrace{W \otimes \dots \otimes W}_{n+k-1}$$

For $\sum_{i_1, \dots, i_n} A_{i_1 i_2 \dots i_n} \in \underbrace{W \otimes \dots \otimes W}_n$ and $\sum_{i_n, \dots, i_{n+k}} B_{i_n i_{n+1} \dots i_{n+k}} \in \underbrace{W \otimes \dots \otimes W}_{k+1}$

Element wise multiplication between two vectors, or matrices, or tensors of the same rank is also an instance of tensor contraction, where one of the arguments of the multiplication is *raised* to a tensor of a higher rank, with the argument in its diagonal and its other entries padded with zero. For an instance of this see Kartsaklis (2016) where coordination is treated in a DisCoCat model, therein the author shows how the linear algebraic closed form of element wise multiplication arises as a result of a tensor contraction.

3.2.2 Lexical Substitution

To obtain a concrete model for a phrase, we need to replace the abstract meaning term of a proof by a concrete tensor mapping. Since we map lambda terms to lambda terms, we only need to specify how constants c are mapped to tensors. This will automatically induce a type-respecting term homomorphism $\mathcal{H}$. A general map that sends constants to a contraction friendly model is presented in Table 1.

The different composition operators of Table 1 seem to be different: we have matrix multiplication for adjectival phrases, intransitive sentences and verb phrases, cube multiplication for transitive sentences, and pointwise multiplication for the conjunctive coordination.

Using Table 1, we can translate the proof term of Fig. 2 as follows:

$$(\text{and } ((\text{dt drinks}) \text{ bob}))(\text{drinks alice}) : s$$

and substitute the concrete terms to get the following β -reduced version:

$$\rightarrow_{\beta} (\mathbf{drinks} \times_1 \mathbf{alice}) \odot (\mathbf{drinks} \times_1 \mathbf{bob})$$

Table 2 Translation that sends abstract terms to a multiplicative-additive model

w	$\sigma(w)$	$\mathcal{H}(w)$	$\mathcal{T}(w)$
cn	n	cn	V
adj	np/n	$\lambda v.(\mathbf{adj} + v)$	VV
adv	$(np \backslash s) \backslash (np \backslash s)$	$\lambda m.(\mathbf{adv} + m)$	VV
itv	$np \backslash s$	$\lambda v.(\mathbf{itv} + v)$	VV
tv	$(np \backslash s)/np$	$\lambda uv.(\mathbf{tv} + u + v)$	VVV
coord	$(s \backslash s)/s$	$\lambda P.Q.(P \odot Q)$	VVV

As another alternative, we can instantiate the proof terms in a multiplicative-additive model. This is a model where the sentences are obtained by adding their individual word embeddings and the overall result is obtained by multiplying the two sentence vectors. This model is presented in Table 2, according to which we obtain the following semantics for our example sentence above:

$$\rightarrow_{\beta} (\mathbf{drinks} + \mathbf{alice}) \odot (\mathbf{drinks} + \mathbf{bob})$$

Another alternative is Table 3, which provides the same terms with a Kronecker - based tensor semantics, originally used by Grefenstette and Sadrzadeh (2011a) to model transitive sentences.

We symbolise the semantics of the basic elliptical phrase that comes out of any of these models for our example sentence as follows:

$$M(\mathbf{sub}_1, \mathbf{verb}) \star M(\mathbf{sub}_2, N(\mathbf{verb}))$$

where M is a general term for an intransitive sentence, N is a term that modifies the verb tensor through the auxiliary verb, and $\star$ is an operation that expresses the coordination of the two subclauses. For a transitive sentence version, the above changes to the following:

$$M(\mathbf{subj}_1, \mathbf{verb}, \mathbf{obj}_1) \star M(\mathbf{subj}_2, N(\mathbf{verb}), \mathbf{obj}_2)$$

Such a description is very general, and in fact allows us to derive almost all compositional vector models that have been tested in the literature (see e.g., Milajevs et al. 2014). This flexibility is necessary for ellipsis because it can model the Cartesian behaviour that is unavailable in a categorical modelling of vectors and linear maps. Some models can, however, only be incorporated by changing the lexical formulas associated to the individual words. The proposal of Kartsaklis et al. (2016) is one such example. They use the coordinator to a heavy extent and their typing and vector/tensor assignments result in the following lambda semantics for the phrase ‘‘Alice drinks and Bob does too’’:

$$\mathbf{drinks} \times_1 (\mathbf{alice} \odot \mathbf{bob})$$

The above is obtained by assigning an identity linear map to the auxiliary phrase ‘does too’ and then assigning a complex linear map to the coordinator ‘and’ tailored in a way that it guarantees the derivation of the final meaning. In our framework, we would

Table 3 Translation that sends abstract terms to a Kronecker model. We abuse the notation to denote the element wise multiplication of two matrices with the same symbol, i.e. $\odot$, as the element wise multiplication of two vectors

w	$\sigma(w)$	$\mathcal{H}(w)$	$\mathcal{T}(w)$
cn	n	cn	V
adj	np/n	$\lambda v.(\mathbf{adj} \odot v)$	VV
adv	$(np \setminus s) \setminus (np \setminus s)$	$\lambda m.(\mathbf{adv} \odot m)$	VV
itv	$np \setminus s$	$\lambda v.(\mathbf{itv} \odot v)$	VV
tv	$(np \setminus s)/np$	$\lambda uv.(\mathbf{tv} \odot (v \otimes u))$	VVM
coord	$(s \setminus s)/s$	$\lambda P.\lambda Q.(P \odot Q)$	VVV

need to take a similar approach, and we need to modify M to essentially return the verb-subject pair, N would be the identity, and and has to be defined with the tailored to purpose term below, which takes two pairs of subjects and verbs, but discards one copy of the verb to mimic the model of Kartsaklis et al. (2016):

$$\text{and} \quad \lambda \langle s, v \rangle. \lambda \langle t, w \rangle. v \times_1 (s \odot t)$$

In either case, we can reasonably derive a large class of compositional functions that can be experimented with in a variety of tasks. With these tools in hand, we can give the desired interpretation to elliptical sentences in the next section.

4 Deriving Ellipsis: Strict and Sloppy Readings

In his book Jäger (2006), Jäger describes various applications of his logic **LLC**. With chapter 5 devoted to verb phrase ellipsis, he discusses various examples of general ellipsis: right node raising, gapping, stripping, VP ellipsis, antecedent contained deletion, and sluicing. Using these categories, an account is developed for VP ellipsis and sluicing. This treatment directly carries over to the vectorial setting, with the challenge that we need to think about how to fill in the lexical semantics. We already gave the basic example of an elliptical phrase in Fig. 2. In this section we show how the account of Jäger allows us to give compositional meanings to ellipsis with anaphora, and cascaded ellipsis, contrasting it with the direct categorical approach, which we show in Wijnholds and Sadzadeh (2018) to be unsuitable for these cases.

4.1 Ellipsis with Anaphora

The interaction of ellipsis with anaphora leads to strict and sloppy readings, as already demonstrated in Sect. 2. We repeat the example here and give the separate derivations:

- a “Gary loves his code and Bob does too” (ambiguous)
- b “Gary loves Gary’s code and Bob loves Bob’s code” (sloppy) (3)
- c “Gary loves Gary’s code and Bob loves Gary’s code” (strict)

Fig. 3 Sloppy interpretation for “Gary loves his code and Bob does-too”: Gary loves Gary’s code and Bob loves Bob’s code

Fig. 4 Strict interpretation for “Gary loves his code and Bob does-too”: Gary loves Gary’s code and Bob loves Gary’s code

The lexical assignment of type $np|(np/n)$ to the possessive pronoun *his* renders it an unbound anaphora, looking for a preceding noun phrase to bind to it. Similarly, the type assignment $(np\s)|(np\s)$ registers ‘does too’ as the ellipsis marker that needs to be bound by a preceding verb phrase. The derivations of the strict (Fig. 4) and the sloppy (Fig. 3) readings essentially differ in their *order of binding*: by binding ‘Gary’ to the possessive pronoun and then binding the resulting verb phrase for ‘loves his code’ to the ellipsis marker, we obtain the strict reading, whereas binding the verb phrase with the unbound possessive pronoun and subsequently binding the two copies of the pronoun differently, we get the sloppy reading. The flexibility of Jäger’s approach is illustrated by the fact that one can ultimately abstract over the binding noun phrase to obtain a third reading, which would derive the type $np\s$, since that pronoun was left unbound.

If we assume a tensor-based compositional model that uses tensor contraction to obtain the meaning of a sentence, we get the two different meanings for the strict and sloppy readings as follows:

1. $((\text{loves} \times_2 (\text{gary} \odot \text{code})) \times_1 \text{gary}) \odot ((\text{loves} \times_2 (\text{gary} \odot \text{code})) \times_1 \text{bob})$ (strict)
2. $((\text{loves} \times_2 (\text{gary} \odot \text{code})) \times_1 \text{gary}) \odot ((\text{loves} \times_2 (\text{bob} \odot \text{code})) \times_1 \text{bob})$ (sloppy)

4.2 Cascaded Ellipsis

Jäger also describes the phenomenon of *cascaded ellipsis*, in which an ellipsis contains an elided verb phrase within itself, as in “Gary revised his code before the student did, and Bob did too”. In this case there are three derivations possible (although even more readings could be found):

1. Gary revised Gary’s code before the student revised Gary’s code, and Bob revised Gary’s code before the student revised Gary’s code.
and
(before (revise ((his gary) code) student) (revise ((his gary) code) gary))
(before (revise ((his gary) code) student) (revise ((his gary) code) bob))
2. Gary revised Gary’s code before the student revised Gary’s code, and Bob revised Bob’s code before the student revised Bob’s code.
and
(before (revise ((his gary) code) student) (revise ((his gary) code) gary))
(before (revise ((his bob) code) student) (revise ((his bob) code) bob))
3. Gary revised Gary’s code before the student revised the student’s code, and Bob revised Bob’s code before the student revised the student’s code.
and
(before (revise ((his student) code) student) (revise ((his gary) code) gary))
(before (revise ((his student) code) student) (revise ((his bob) code) bob))

A tensor-based model would assign three meanings appropriately. For example, the first subclause of 1 would give the following

$$(\text{gary} \times_1 \text{revise} \times_2 (\text{gary} \odot \text{code})) \star (\text{student} \times_1 \text{revise} \times_2 (\text{gary} \odot \text{code}))$$

where $\star$ interprets the function word ‘before’.

In the next section we carry out experimental evaluation of the framework developed so far. We start out with a toy experiment and then perform a large-scale experiment on verb phrase-elliptical sentences. We do not cover the more complex cases of ellipsis that involve ambiguities: setting up experiments for those cases is a task on its own and requires more investigation.

5 Experimental Evaluation

To evaluate the framework we have developed so far, we carry out an experiment involving verb disambiguation. This kind of task was initiated in the work of Mitchell and Lapata (2008) and Mitchell and Lapata (2010) in order to evaluate the compositional vectors of intransitive sentences and verb phrases. These have been extended to transitive sentences (Grefenstette and Sadrzadeh 2011a; Kartsaklis and Sadrzadeh 2013). Here, we introduce the general idea behind the verb disambiguation task and how it is solved with compositional distributional models, before proceeding to an illustrative toy experiment and a large scale experiment.

A distributional model on the word level is considered successful if it optimises the similarity between words. Whenever two words w_1 and w_2 are considered similar, the associated vectors $\vec{w}_1$ and $\vec{w}_2$ ought to be similar as well. Similarity judgments between words are obtained by asking human judges, whereas the customary way of measuring similarity between vectors is given by the cosine of the angle between vectors (cosine similarity):

$$\cos(\vec{v}, \vec{w}) = \frac{\vec{v} \cdot \vec{w}}{|\vec{v}| |\vec{w}|}$$

where $\cdot$ denotes the dot product and $|\cdot|$ denotes the magnitude of a vector.

Compositional tasks follow the same pattern, but now one is also interested in (a) how context affects similarity judgments and (b) how word representations are to be composed to give a sentence vector. The idea behind the verb disambiguation tasks (Mitchell and Lapata 2008; Grefenstette and Sadrzadeh 2011a; Kartsaklis and Sadrzadeh 2013) is that sentences containing an ambiguous verb can be disambiguated by context. An example is the verb *meet* which can mean *visit* or *satisfy* (a requirement). In the sentence *Students meet teachers*, *meet* means *visit*, whereas in the sentence *Houses meet standard*, it means *satisfy*. What makes this idea suitable for compositional distributional semantics is that we can use the vectors of these sentences to disambiguate the verb. This is detailed below.

Suppose we have a verb V that is ambiguous between two different meanings V_1 and V_2 , we refer to V_1 and V_2 as the two landmark meanings of V . We position V in a sentence $Sbj\ V\ Obj$ in which only one of the meanings of the verbs makes sense. Suppose that meaning is V_1 , so the sentences $Sbj\ V\ Obj$ and $Sbj\ V_1\ Obj$ make sense while the sentence $Sbj\ V_2\ Obj$ does not. Then the cosine similarity between the vectors for the first two sentences ought to be high, but between those for the first and the third sentence it ought to be low. So the hypothesis that is tested is that this disambiguation by context manifests when we compute vectors for the meanings of these sentences. In technical terms, we wish the distance between the meaning vector of $Sbj\ V_1\ Obj$ and that of $Sbj\ V\ Obj$ to be smaller than the distance between the vector of $Sbj\ V_2\ Obj$ and that of $Sbj\ V\ Obj$:

$$\cos(\overrightarrow{Sbj\ V\ Obj}, \overrightarrow{Sbj\ V_1\ Obj}) \geq \cos(\overrightarrow{Sbj\ V\ Obj}, \overrightarrow{Sbj\ V_2\ Obj})$$

This hypothesis forms the basis for our verb disambiguation task. Each of the datasets created for verb disambiguation (Mitchell and Lapata 2008; Grefenstette

and Sadrzadeh 2011a; Kartsaklis and Sadrzadeh 2013) contains a balanced number of subjects, or subject-object combinations for several verbs and two landmark interpretations. That is, for a verb V ambiguous between V_1 and V_2 there will be roughly an equal number of contexts that push the meaning of V to V_1 and to V_2 . Moreover, these datasets contain similarity judgments that allow us to not just classify the most likely interpretation of a given verb, but to compute the correlation between a model's prediction and the human judgments, to see how well the model aligns with humans.

The basic such models for composing word vectors to sentence vectors are the additive and multiplicative models that, for any sentence, simply add or multiply the vectors for the words in the sentence. For intransitive sentences of the form $Sbj\ V$, we would get respectively

$$\overrightarrow{Sbj\ V} = \overrightarrow{Sbj} + \overrightarrow{V} \quad \overrightarrow{Sbj\ V} = \overrightarrow{Sbj} \odot \overrightarrow{V}$$

For the transitive case, of the form $Sbj\ V\ Obj$, we additionally consider the Kronecker model used of Grefenstette and Sadrzadeh (2011b), which assigns to a sentence *subj verb obj* the following formula:

$$(\overrightarrow{verb} \otimes \overrightarrow{verb}) \odot (\overrightarrow{subj} \otimes \overrightarrow{obj})$$

In this model, note that the resulting representation is now a matrix rather than a vector.

Here, we extend the experimental setting to elliptical sentences, our hypothesis is twofold: on the one hand, an elliptical phrase will have more content that adds to the context of the verb to be disambiguated, allowing us to disambiguate more effectively. On the other hand, we test the disambiguating effect of resolving the ellipsis. Going to an elliptical setting allows us to define several more composition models based on the additive/multiplicative and Kronecker models: for a transitive sentence $Sbj\ V\ Obj$ extended to the elliptical setting $Sbj\ V\ Obj\ and\ Sbj'\ does\ too$, we can again consider the additive and multiplicative models:

$$\begin{aligned} \overrightarrow{Sbj\ V\ Obj\ and\ Sbj'\ does\ too} &= \overrightarrow{Sbj} + \overrightarrow{V} + \overrightarrow{Obj} + \overrightarrow{and} + \overrightarrow{Sbj'} + \overrightarrow{does} + \overrightarrow{too} \\ \overrightarrow{Sbj\ V\ Obj\ and\ Sbj'\ does\ too} &= \overrightarrow{Sbj} \odot \overrightarrow{V} \odot \overrightarrow{Obj} \odot \overrightarrow{and} \odot \overrightarrow{Sbj'} \odot \overrightarrow{does} \odot \overrightarrow{too} \end{aligned}$$

In addition, we can now consider combinations of models on the resolved elliptical phrases, following the pattern of Sect. 3. For an intransitive as well as a transitive sentence extended to an elliptical setting, its resolved version combines the two implicit subclauses by an operation. Hence, we can use one of the models outlined above on the subclauses, and then choose an operation to combine them. This leads, for the intransitive case, to the following four models:

Model	Formula
Multiplicative $\odot$	$\overrightarrow{subj} \odot \overrightarrow{verb} \odot \overrightarrow{subj}^* \odot \overrightarrow{verb}$
Multiplicative +	$(\overrightarrow{subj} \odot \overrightarrow{verb}) + (\overrightarrow{subj}^* \odot \overrightarrow{verb})$
Additive $\odot$	$(\overrightarrow{subj} + \overrightarrow{verb}) \odot (\overrightarrow{subj}^* + \overrightarrow{verb})$
Additive +	$\overrightarrow{subj} + \overrightarrow{verb} + \overrightarrow{subj}^* + \overrightarrow{verb}$

For the transitive case, we additionally get the Kronecker + and Kronecker $\odot$ models, given by either summing or multiplying the two Kronecker model matrices of the subclauses:

Model	Formula
Kronecker +	$(\overrightarrow{verb} \otimes \overrightarrow{verb}) \odot (\overrightarrow{subj} \otimes \overrightarrow{obj}) + (\overrightarrow{verb} \otimes \overrightarrow{verb}) \odot (\overrightarrow{subj}' \otimes \overrightarrow{obj})$
Kronecker $\odot$	$(\overrightarrow{verb} \otimes \overrightarrow{verb}) \odot (\overrightarrow{subj} \otimes \overrightarrow{obj}) \odot (\overrightarrow{verb} \otimes \overrightarrow{verb}) \odot (\overrightarrow{subj}' \otimes \overrightarrow{obj})$

5.1 A Toy Experiment

In order to demonstrate the effect of vectors and distances in this task, we provide a hypothetical though intuitive example. Consider the sentence “the man runs”, which is ambiguous between “the man races” and “the man stands (for election)”. The sentence itself does not have enough context to help disambiguate the verb, but if we add a case of ellipsis such as “the man runs, the dog too” to it, the ambiguity will be resolved. Another example, this time transitive, is the sentence “the man draws the sword”, which is ambiguous between “the man pulls the sword” and “the man depicts the sword”. Again, the current sentence in which the ambiguous verb occurs may not easily disambiguate it, but after adding the extra context “the soldier does too”, the disambiguating effect of the context is much stronger.

Consider the following vector space built from raw co-occurrence counts of several nouns and verbs with respect to a set of context words. The co-occurrence matrix is given in Table 4; each row of the table represents a word embedding.

We work out the cosine similarity scores between vector representations of a sentence with an ambiguous verb and its two landmark interpretations, following the models outlined above, on the concrete sentence “the man runs” with the extension of ‘governor’ and ‘athlete’ respectively. The idea is that the representation of “the man runs and governor does too” will be closer to that of “the man stands and governor does too”, whereas the representation of “the man runs and athlete does too” will be closer to that of “the man races and athlete does too”. The cosine similarity scores for each model are presented in Table 5. The original representation of “the man runs” is more similar to that of “the man races” by a difference of 0.10. However, for all models except the fully additive model, we see that adding the extra subject increases the difference between similarity scores, thereby making it easier to distinguish the

Table 4 (Hypothetical) co-occurrence counts for several nouns and verbs

	Human	Painting	Army	Weapon	Marathon	Election
man	2	3	4	2	4	4
painter	3	8	1	3	1	1
warrior	4	1	2	9	1	0
sword	2	3	9	2	0	0
picture	1	20	0	1	1	1
governor	7	1	1	3	1	9
athlete	6	2	0	1	9	1
draw	4	10	9	11	2	3
pull	7	2	10	15	1	1
depict	3	15	2	2	1	2
run	4	0	2	1	8	7
race	8	0	0	3	10	3
stand	5	1	0	1	2	11

Table 5 Cosine similarity scores between representations involving the intransitive verb ‘run’. Column **race**: the representation of the corresponding row sentence but with ‘race’ instead of ‘run’, similarly for **stand**

	Multiplicative $\odot$		Multiplicative $+$		Additive $\odot$		Additive $+$	
	race	stand	race	stand	race	stand	race	stand
Man run	.88	.78	.88	.78	.94	.92	.94	.92
Man run, governor does too	.47	.99	.80	.94	.82	.89	.95	.93
Man run, athlete does too	.99	.36	.96	.71	.94	.71	.95	.92

Bold represents the results for the highest performing models

correct interpretation. The most discriminative model is the fully multiplicative one, which treats the conjunctive coordinator as multiplication.

For the transitive case, we compare the sentences “the man draws the sword” and “the man draws the picture” with their landmark interpretations in which the verb ‘draw’ is replaced by either ‘pull’ or ‘depict’. All of these are extended with the contexts ‘warrior’ and ‘painter’, and we compute the result of four of the mixed transitive models outlined above for the elliptical case: two are the same additive models that just sum all the vectors in a (sub)clause and either sum or multiply vectors for the subclauses for the elliptical variant, and two models use the Kronecker representation detailed above. The concrete cosine similarity scores are displayed in Table 6.

In this case, neither of the additive models seem to be effective: for the original phrases they already give very high similarity scores, and those do not change greatly after adding the extra context. For the Kronecker models, we see that the best discriminatory model is the one that multiplies the vectors for the subclauses: in both original transitive phrases the interpretation ‘pull’ is more similar than ‘depict’, but adding the context improves the disambiguation results. For the first phrase, where a

Table 6 Cosine similarity scores between sentence representations using several models. Column **pull**: the representation of the corresponding row sentence but with ‘pull’ instead of ‘draw’, similarly for **depict**

	Kronecker $\ominus$		Kronecker $+$		Additive $\ominus$		Additive $+$	
	pull	depict	pull	depict	pull	depict	pull	depict
Man draw sword	.83	.50	.83	.50	.96	.93	.96	.93
Man draw sword, warrior does too	.98	.07	.92	.44	.94	.76	.96	.93
Man draw sword, painter does too	.37	.28	.69	.59	.89	.80	.96	.93
Man draw picture	.82	.74	.82	.74	.97	.95	.97	.95
Man draw picture, warrior does too	.98	.25	.91	.65	.92	.97	.97	.95
Man draw picture, painter does too	.37	.95	.68	.88	.96	.98	.97	.96

Bold represents the results for the highest performing models

sword is drawn, the addition of ‘warrior’ greatly improves the similarity with ‘pull’ and accordingly decreases the similarity with ‘depict’, though for the addition of ‘painter’ this is not the case. The representation for ‘painter’ is in itself already closer to that of ‘depict’ (cosine similarity of 0.97) than it is to that of ‘pull’ (cosine similarity of 0.52), so adding ‘painter’ to the sentence makes it harder to be certain about ‘pull’ as a likely interpretation of ‘draw’. We see in fact that the difference between the two sentence interpretations has become smaller.

For the second phrase, in which a picture is drawn, the original ambiguity is bigger, but adding the context provides us with the appropriate disambiguating scores. As with the first phrase, we also experience the difficulty in disambiguation: a human may deem “man draw picture, warrior does too” to be more similar to “man depict picture, warrior does too” since pulling a picture is not a very sensible action. However, because the vector for ‘warrior’ is closer to that for ‘pull’ (cosine similarity of 0.94) than it is to ‘depict’ (cosine similarity of 0.31) the model will favour the interpretation in which the picture is pulled.

5.2 Large Scale Evaluation

In addition to a hypothetical toy example, we experimented with our models on a large scale dataset, obtained by extending the disambiguation dataset of Mitchell and Lapata (2008), which we will refer to as the The ML2008 dataset dataset. The ML2008 is an instance of the verb disambiguation task that we have been discussing so far, and contains 120 pairs of sentences: for each of 15 verb triples (V , V_1 , V_2), where verb V is ambiguous between interpretation V_2 and V_3 , four different context subjects were added, and the so constructed sentence pairs $SbjV$, $SbjV_1$ and $SbjV$, $SbjV_2$ were annotated for similarity by humans. For each of two sentence pairs, the interpretation that was assumed more likely was labelled HIGH before collecting annotations, and the other one was labelled LOW; this was done both for verification purposes as well as randomisation of the presentation of the sentence pairs to human judges. The subjects that were added would be mixed: some would cause the verb to tend to one interpretation, others cause the verb to be interpreted with the second meaning.

For example, the dataset contains the pairs

Landmark	High	Low
Export boom	Export prosper	Export thunder
Gun boom	Gun thunder	Gun prosper

To extend such pairs to an elliptical setting, we chose a second subject for each sentence, as follows: for a given subject/verb combination and its two interpretations, we chose a new subject that occurred frequently in a corpus,² but significantly more frequently with the more likely unambiguous verb (the one marked HIGH). For example, the word “economy” occurs with “boom” but it occurs significantly more often with “prosper” than it does with “thunder”. And similarly, “cannon” occurs with “boom” and “thunder” but not so often with “prosper”. We then format the pairs from the ML2008 dataset using the new subject and the elliptical setting. For the examples above, we then got

Landmark	Export boom and economy does too	Gun boom and cannon does too
HIGH	Export prosper and economy does too	Gun thunder and cannon does too
LOW	Export thunder and economy does too	Gun prosper and cannon does too

In total, we added two new subjects to each sentence pair, producing a dataset of 240 entries. We used the human similarity judgments of the original ML2008 dataset to see whether the addition of disambiguating context, combined with our ellipsis model, will be able to better distinguish verb meaning. As explained in the start of this section, we use several different concrete models to compute the representation of the sentences in the dataset, and compute the cosine similarity between sentences in a pair; the predicted judgments are then evaluated by computing the (linear) degree of correlation with human similarity judgments, using the standard Spearman ρ measure (Tables 7, 8).

We used two different instantiations of a vector space model: the first is a 300-dimensional model trained on the Google News corpus, taken from the popular and widely used `word2vec` package³, which is based on the Skipgram model with negative sampling of Mikolov et al. (2013). This model is known to lead to high-quality dense vector embeddings of words. The second space we used is a custom trained 2000-dimensional vector space, trained on the combined UKWaC and Wackypedia corpus, using a context window of 5, and Positive Pointwise Mutual Information as a normalisation scheme on the raw co-occurrence counts. The vectors of this space do not involve any dimensionality reduction techniques, making the vectors relatively sparse compared to those in the `word2vec` vector space.

² In our case, this was the combined UKWaC and Wackypedia corpus, available at www.wacky.sslmit.unibo.it.

³ <https://code.google.com/p/word2vec/>.

Table 7 Spearman ρ correlation scores on the **ML2008** dataset

ML2008	word2vec	Count Based
Verb Only Vector	0.274	0.078
Verb Only Tensor	0.060	0.108
Additive	0.278	0.081
Multiplicative	0.229	0.177

Bold represents the results for the highest performing models

Table 8 Spearman ρ correlation scores on the extended **ML2008** dataset

MLELLDIS	word2vec	Count Based
Verb Only Vector	0.274	0.078
Verb Only Tensor	0.060	0.108
Additive	0.292	0.040
Multiplicative	0.068	0.206
Multiplicative $\odot$	0.213	0.391
Multiplicative +	0.294	0.179
Additive $\odot$	0.229	0.172
Additive +	0.298	0.078

Bold represents the results for the highest performing models

For the original dataset, we compare a non-compositional baseline, in which just the vector or matrix for the verb is compared, and additive/multiplicative models, and get the results in Table 7.

These results are higher than found in the literature: the original evaluation of Mitchell and Lapata (2008) achieved a highest correlation score of 0.19, and the regression model of Grefenstette et al. (2013) achieves a top correlation score of 0.23. These scores are surpassed already by the non-compositional baseline on the word2vec space here. Although the highest scores are indeed obtained using a compositional model, note that the correlation for the word2vec model doesn't increase substantially. In the count based space we do see a bump in the correlation when using a compositional model, but here the baseline correlation isn't that high to start with. The situation is better for the extended dataset. There, we compare the same four models against four combined models, which combine and additive with a multiplicative model, after resolving the ellipsis. The results are in Table 8.

Our first observation is that the naive additive and multiplicative models already do better than the non-compositional baseline, save for the additive model on the count based space. Secondly, even better results are obtained by applying a non-linear compositional model, i.e. a model that actually resolves the ellipsis and copies the representation of the verb. For the case of the word2vec space the best performing model is the fully additive model that adds together all the vectors to give the result $\vec{subj} + \vec{verb} + \vec{subj}^* + \vec{verb}$. For the count based space, it is the exact opposite: the fully multiplicative model achieves the best overall score of 0.391 with the representation $\vec{subj} \odot \vec{verb} \odot \vec{subj}^* \odot \vec{verb}$.

That the `word2vec` vectors work well under addition but not under multiplication, whereas the count based vectors work well under multiplication but not under addition, we attribute to the difference in sparsity of the representations: since `word2vec` vectors are very dense representations, multiplying them will not have a very strong effect on the resulting representation, whereas adding them will have a greater net effect on the final result. In contrast, multiplying two sparse vectors will eliminate a lot of information, since the entries that are zero in one of the vectors leads to a zero entry in the final vector. In other words, the final representation will be incredibly specific, allowing for better disambiguation. Addition on sparse vectors however, will simply generate vectors that are very unspecific and are thus not very helpful for disambiguation.

Overall, we see that the presented results are in favour of non-linear compositional models, showing the importance of ellipsis resolution for distributional sentence representations.

6 Conclusion, Further Work

In this paper we incorporated a proper notion of copying into a compositional distributional model of meaning to deal with VP ellipsis with anaphora. By decomposing the DisCoCat architecture into a two step interpretation process, we were able to combine the flexibility of the Cartesian structure of the non-linear simply typed lambda calculus, with a vector based representation of word meaning. We presented a vector-based analysis of VP ellipsis with anaphora and showed how the elliptical phrases get assigned the same meaning as their resolved variants. We also carried out a large scale similarity experiment, showing that verb disambiguation becomes easier after ellipsis resolution.

By giving up a direct categorical translation from a typological grammar to vector spaces, we gain the expressiveness of the lambda calculus, which allows one to interpret the grammatical derivations in various different concrete compositional models of meaning. We showed that previous DisCoCat work on resolving ellipsis using coordination and Frobenius algebras (Kartsaklis et al. 2016) can only be obtained in an ad hoc fashion. For future work we intend to compare the two approaches from an experimental point of view.

A second challenge that we would like to address in the future involves dealing with derivational ambiguities in a vectorial setting. These ambiguities were exemplified in this paper by the strict and sloppy readings of elliptical phrases involving anaphora, and cascaded ellipsis. In order to experiment with the vectorial models of these cases, an appropriate task should be defined and experiments should determine which distributional reading can be empirically validated.

Finally, in previous work (Wijnholds and Sadrzadeh 2018), we showed how to resolve ellipsis in a modal Lambek Calculus which has a controlled form of contraction for formulae marked with the modality. Our work is very similar to an earlier proposal of Jäger presented in Jäger (1998). The controlled contraction rule that we use is as follows

$$\frac{f : \Diamond A \otimes A \rightarrow B}{C(f) : A \rightarrow B}$$

The $\diamond$ modality has a few other rules for controlled associativity and movement. The semantics of this rule is, however, simply defined as $C := \lambda f x. f \langle x, x \rangle$. Trying to find a vector operation (either in a linear setting using a biproduct operation or by moving to a non-linear setting) and obtaining a direct categorical semantics is work in progress. The challenge is that the interpretations of the similar ! modality of Linear Logic, e.g. in a linguistic setting by Morrill in Morrill (2017) or in a computational setting by Abramsky in Abramsky (1993)) would not work in a vector space setting. The Frobenius algebraic copying operation, with which we worked in Wijnholds and Sadrzadeh (2018), is one of the few options available, and we have shown that it does not work when it comes to distinguishing the sloppy versus strict reading of the ambiguous elliptical cases.

Acknowledgements The authors gratefully acknowledge support from the Royal Society International Exchange Award IE161631—Dynamic Vector Semantics for Lambda Calculus Models of Natural Language—and a QMUL Principal Studentship.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Abramsky, S. (1993). Computational interpretations of linear logic. *Theoretical Computer Science*, 111(1), 3–57.
- Abramsky, S. (2009). No-cloning in categorical quantum mechanics. *Semantic Techniques in Quantum Computation*, 1–28.
- Baroni, M., & Zamparelli, R. (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). Association for Computational Linguistics.
- Coecke, B., Grefenstette, E., & Sadrzadeh, M. (2013). Lambek vs. Lambek: Functorial vector space semantics and string diagrams for lambek calculus. *Annals of Pure and Applied Logic*, 164(11), 1079–1100.
- Coecke, B., & Paquette, E. (2008). Introducing categories to the practicing physicist. In B. Coecke (Ed.), *New structures for physics, lecture notes in physics* (pp. 167–271). Berlin: Springer.
- Coecke, B., Pavlovic, D., & Vicary, J. (2008). A new description of orthogonal bases. *Mathematical Structures in Computer Science*, 1, 269–272.
- Coecke, B., Sadrzadeh, M., & Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. arXiv preprint [arXiv:1003.4394](https://arxiv.org/abs/1003.4394).
- Dalrymple, M., Shieber, S. M., & Pereira, F. C. (1991). Ellipsis and higher-order unification. *Linguistics and Philosophy*, 14(4), 399–452.
- de Groote, P. (2001). Towards abstract categorial grammars. In *Proceedings of the 39th Annual Meeting on Association for Computational Linguistics, ACL '01* (pp. 252–259). Association for Computational Linguistics, Stroudsburg, PA, USA.
- Firth, J. R. (1957). A synopsis of linguistic theory, 1930–1955. *Studies in Linguistic Analysis 1952–59*, 1–32.
- Grefenstette, E., Dinu, G., Zhang, Y. Z., Sadrzadeh, M., & Baroni, M. (2013). Multi-step regression learning for compositional distributional semantics. arXiv preprint [arXiv:1301.6939](https://arxiv.org/abs/1301.6939).
- Grefenstette, E., & Sadrzadeh, M. (2011a). Experimental support for a categorial compositional distributional model of meaning. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 1394–1404). Association for Computational Linguistics.
- Grefenstette, E., & Sadrzadeh, M. (2011b). Experimenting with transitive verbs in a discocat. In *Proceedings of the GEMS 2011 workshop on geometrical models of natural language semantics* (pp. 62–66). Association for Computational Linguistics.

- Hedges, J., & Sadrzadeh, M. (2016). A generalised quantifier theory of natural language in categorical compositional distributional semantics with bialgebras. arXiv preprint [arXiv:1602.01635](https://arxiv.org/abs/1602.01635).
- Hendriks, P. (1995). *Comparatives and Categorical Grammar*. Groningen dissertations in linguistics. Grodil. <https://books.google.co.uk/books?id=FeNUuAAACAAJ>.
- Jacobs, B. (2011). Bases as coalgebras. In *International conference on algebra and coalgebra in computer science*, (pp. 237–252). Berlin: Springer.
- Jäger, G. (1998). A multi-modal analysis of anaphora and ellipsis. *University of Pennsylvania Working Papers in Linguistics*, 5(2), 2.
- Jäger, G. (2006). *Anaphora and type logical grammar* (Vol. 24). Berlin: Springer.
- Kartsaklis, D. (2016). Coordination in categorical compositional distributional semantics. arXiv preprint [arXiv:1606.01515](https://arxiv.org/abs/1606.01515).
- Kartsaklis, D., Purver, M., & Sadrzadeh, M. (2016). Verb phrase ellipsis using frobenius algebras in categorical compositional distributional semantics. In *European summer school on logic, language and information: DSALT workshop*.
- Kartsaklis, D., & Sadrzadeh, M. (2013). Prior disambiguation of word tensors for constructing sentence vectors. In *Proceedings of the 2013 conference on empirical methods in natural language processing* (pp. 1590–1601).
- Kempson, R., Cann, R., Eshghi, A., Gregoromichelaki, E., & Purver, M. (2015). Ellipsis. In S. Lappin & C. Fox (Eds.), *Handbook of contemporary semantic theory*, chap. 4 (2nd ed.). Hoboken: Wiley.
- Kruszewski, G., Paperno, D., Bernardi, R., & Baroni, M. (2016). There is no logical negation here, but there are alternatives: Modeling conversational negation with distributional semantics. *Computational Linguistics*, 42(4), 637–660.
- Kubota, Y., & Levine, R. (2017). Pseudogapping as pseudo-vp-ellipsis. *Linguistic Inquiry*, 48(2), 213–257.
- Lambek, J. (1997). Type grammar revisited. In *International conference on logical aspects of computational linguistics* (pp. 1–27). Berlin: Springer.
- Lenci, A. (2008). Distributional semantics in linguistic and cognitive research. *Italian Journal of Linguistics*, 20(1), 1–31.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers*, 28(2), 203–208.
- Maillard, J., Clark, S., & Grefenstette, E. (2014). A type-driven tensor-based semantics for CCG. In *Proceedings of the EACL 2014 type theory and natural language semantics workshop* (pp. 46–54).
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems* (pp. 3111–3119).
- Milajevs, D., Kartsaklis, D., Sadrzadeh, M., & Purver, M. (2014). Evaluating neural word representations in tensor-based compositional settings. arXiv preprint [arXiv:1408.6179](https://arxiv.org/abs/1408.6179).
- Mitchell, J., & Lapata, M. (2008). Vector-based models of semantic composition. In *Proceedings of ACL-08: HLT* (pp. 236–244).
- Mitchell, J., & Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive Science*, 34(8), 1388–1429.
- Montague, R. (1970). English as a formal language. In B. Visentini (Ed.), *Linguaggi nella società e nella tecnica* (pp. 188–221). Ivrea: Edizioni di Comunità.
- Moortgat, M. (1997). Chapter 2—categorical type logics. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 93–177). Amsterdam: Elsevier. <https://doi.org/10.1016/B978-044481714-3/50005-9>.
- Moortgat, M., & Wijnholds, G. (2017). Lexical and derivational meaning in vector-based models of relativisation. In *Proceedings of the 21st Amsterdam colloquium* (pp. 55–64).
- Morrill, G. (2017). Grammar logicised: Relativisation. *Linguistics and Philosophy*, 40(2), 119–163.
- Morrill, G., & Merenciano Saladrigas, J. M. (1996). Generalising discontinuity. *Traitement Automatique des Langues*, 27, 119–143.
- Morrill, G., & Valentin, O. (2010). On calculus of displacement. In *Proceedings of the 10th international workshop on tree adjoining grammars and related formalisms* (pp. 45–52).
- Muskens, R. (2003). Language, lambdas, and logic. *Resource-sensitivity, binding and anaphora* (pp. 23–54). Dordrecht: Springer.
- Muskens, R., & Sadrzadeh, M. (2016). Context update for lambdas and vectors. In M. Amblard, P. de Groote, S. Pogodalla, & C. Retoré (Eds.), *Logical aspects of computational linguistics. Celebrating 20 years of LACL (1996–2016)* (pp. 247–254). Berlin: Springer.

- Muskens, R., & Sadrzadeh, M. (2019). Static and dynamic vector semantics for lambda calculus models of natural language. *Journal of Language Modelling*, 6(2), 319–351.
- Sadrzadeh, M. (2016). Quantifier scope in categorical compositional distributional semantics. arXiv preprint [arXiv:1608.01404](https://arxiv.org/abs/1608.01404).
- Sadrzadeh, M., Clark, S., & Coecke, B. (2013). The Frobenius anatomy of word meanings I: Subject and object relative pronouns. *Journal of Logic and Computation*, 23(6), 1293–1317.
- Sadrzadeh, M., Clark, S., & Coecke, B. (2014). The Frobenius anatomy of word meanings II: Possessive relative pronouns. *Journal of Logic and Computation*, 26, 785–815.
- Schütze, H. (1998). Automatic word sense discrimination. *Computational Linguistics*, 24(1), 97–123.
- Wijnholds, G. (2014). Categorical foundations for extended compositional distributional models of meaning. M.Sc. thesis . <https://www.ilic.uva.nl/Research/Publications/Reports/MoL-2014-22.text.pdf>.
- Wijnholds, G., & Sadrzadeh, M. (2018). Classical copying versus quantum entanglement in natural language: The case of VP-ellipsis. In *EPTCS proceedings of the second workshop on compositional approaches for physics, NLP, and social sciences (CAPNS)*.
- Wijnholds, G. (2019). A proof-theoretic approach to scope ambiguity in compositional vector space models. *Journal of Language Modelling*, 6(2), 261–286.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.