

# Towards Compositional Distributional Discourse Analysis

Bob Coecke, Giovanni de Felice, Dan Marsden & Alexis Toumi

{firstname.lastname}@cs.ox.ac.uk

Department of Computer Science, University of Oxford

Categorical compositional distributional semantics provide a method to derive the meaning of a sentence from the meaning of its individual words: the grammatical reduction of a sentence automatically induces a linear map for composing the word vectors obtained from distributional semantics. In this paper, we extend this passage from word-to-sentence to sentence-to-discourse composition. To achieve this we introduce a notion of basic anaphoric discourses as a mid-level representation between natural language discourse formalised in terms of basic *discourse representation structures* (DRS); and knowledge base queries over the Semantic Web as described by *basic graph patterns* in the Resource Description Framework (RDF). This provides a high-level specification for compositional algorithms for *question answering* and *anaphora resolution*, and allows us to give a picture of *natural language understanding* as a process involving both statistical and logical resources.

## Introduction

In the last couple of decades, the traditional *symbolic* approach to AI and cognitive science — which aims at characterising human intelligence in terms of abstract logical processes — has been challenged by so-called *connectionist* AI: the study of the human brain as a complex network of basic processing units [18]. When it comes to human language, the same divide manifests itself as the opposition between two principles, which in turn induce two distinct approaches to Natural Language Processing (*NLP*). On one hand Frege’s principle of *compositionality* asserts that the meaning of a complex expression is a function of its sub-expressions, and the way in which they are composed — *distributionality* on the other hand can be summed up in Firth’s maxim “You shall know a word by the company it keeps”. Once implemented in terms of concrete algorithms we have expert systems driven by formal logical rules on one end, artificial neural networks and machine learning on the other.

Categorical Compositional Distributional (*DisCoCat*) models, first introduced in [4], aim at getting the best of both worlds: the string diagrams notation borrowed from category theory allows to manipulate the grammatical reductions as linear maps, and compute graphically the semantics of a sentence as the composition of the vectors which we obtain from the distributional semantics of its constituent words.

In this paper, we introduce *basic anaphoric discourses* as mid-level representations between natural language discourse on one end — formalised in terms of basic *discourse representation structures* (DRS) [2]; and knowledge queries over the Semantic Web on the other — given by *basic graph patterns* in the Resource Description Framework (RDF) [19]. We construct discourses as formal diagrams of real-valued matrices and we then use these diagrams to give abstract reformulations of NLP problems: probabilistic *anaphora resolution* and *question answering*.

In the first two sections we introduce the notation used in the rest of the paper, then give a brief summary of how DisCoCat models can be used to turn declarative sentences into probabilistic Ask-type queries in the SPARQL protocol and RDF query language. In the third section, we extend this analysis to Select-type queries and show how they can be translated as “who” or “whom” questions in the simple case, as discourses making use of ambiguous pronouns such as “they” or “them” in more complex cases,



A *knowledge graph* (also called an *ontology*) is given by a set of *RDF triples*  $K \subseteq E \times R \times E$  which contains all the true statements that we care about, e.g.  $(Alice, loves, Bob), (Bob, hates, Charles) \in K$ . We can turn any knowledge graph into an effect by summing over the column vectors of its triples:


$$\langle K | = \sum_{(s,v,o) \in K} \langle s | \otimes \langle v | \otimes \langle o | : |E| \otimes |R| \otimes |E| \rightarrow 1$$

Then for any triple  $t \in E \times R \times E$ , if we compose the state  $|t\rangle : 1 \rightarrow |E| \otimes |R| \otimes |E|$  with the effect  $\langle K |$ , the scalar  $\langle K | t \rangle \in S$  we obtain is equal to 1 if and only if  $t \in K$ — we consider this scalar as the outcome of a basic knowledge graph query. In order to study more complex queries, we need a way to wire knowledge graphs together: we do this using the *generalised Kronecker delta* together with the following standard results.

**Definition.** For  $a, b, n \in \mathbb{N}$ , the  $(a, b)$  Kronecker delta over  $n$  — which we depict as an  $a$ -input,  $b$ -output spider — is the matrix:

$$\delta_n^{a,b} = \sum_{i < n} |e_i\rangle^{(\otimes b)} \langle e_i|^{(\otimes a)} : n^a \rightarrow n^b \quad (1)$$


**Theorem 1.** For all  $n \in \mathbb{N}$ , the set of matrices  $\{\delta_n^{a,b}\}_{a,b \in \mathbb{N}}$  obeys the spider fusion equations:



*Proof.* We have:  $(\text{id}_n^{(\otimes b)} \otimes \delta_n^{(k+c),d}) (\delta_n^{a,(b+k)} \otimes \text{id}_n^{(\otimes c)}) = \delta_n^{(a+c),(b+d)}$ , see [6]. Note how the string diagrams make the topological nature of the equation apparent while hiding the bureaucracy of indices.  $\square$


The  $(2, 0)$  and  $(0, 2)$  Kronecker deltas — which we depict as *cups* and *caps* respectively — make  $\text{Mat}(S)$  a *compact closed* category. We recall a couple of useful properties of cups and caps.

**Corollary 1.1.** For all  $n \in \mathbb{N}$ , the effect  $\text{cup}_n = \delta_n^{2,0} : n \otimes n \rightarrow 1$  and the state  $\text{cap}_n = \delta_n^{0,2} : 1 \rightarrow n \otimes n$  obey the snake equations:



*Proof.* We have:  $(\text{id}_n \otimes \text{cup}_n) (\text{cap}_n \otimes \text{id}_n) = \text{id}_n = (\text{cup}_n \otimes \text{id}_n) (\text{id}_n \otimes \text{cap}_n)$  which follows from Theorem 1 and  $\delta_n^{1,1} = \sum_{i < n} |e_i\rangle \langle e_i| = \text{id}_n$ , see [8].  $\square$

**Corollary 1.2.** We can define the transpose of a matrix  $M : m \rightarrow n$  as:




*Proof.* We have  $M^T = (\text{id}_n \otimes \text{cup}_m) (\text{id}_n \otimes M \otimes \text{id}_m) (\text{cap}_n \otimes \text{id}_m) : n \rightarrow m$ , see [8].  $\square$

## 2 Diagrams for Distributional Compositional Semantics

Lambek introduced *pregroup grammars* as a means of encoding the grammatical structure of natural language. In [4, 5], Clark et al. observe that both pregroups and vector spaces carry the same mathematical structure: both are compact closed categories. This allows us to assign vectors to words and automatically combine them together using a linear map induced by the grammatical reduction. In this section, we give a graphical presentation of the concrete implementation described in [11] and reformulate their construction in terms of knowledge graphs.

The first step is to build a *noun space*  $N$  which we will assign to the noun type of our pregroup grammar. This can be done by collecting co-occurrence data over a large natural language corpus: the dimension  $n$  of the noun space can be seen as a hyper-parameter of the model, where in the simple case each dimension corresponds to a “context word”. However, our model is agnostic of the exact method used to construct  $N$ : we will only assume that every entity  $e$  is assigned a vector  $v_e \in N$ , and store them as one matrix  $E = \sum_{e \in E} v_e \langle e |$ . This encoding matrix  $E : |E| \rightarrow n$  gives us a measure of similarity between the entities, indeed whereas we have  $\langle \text{cat} | \text{dog} \rangle = 0$  now we expect the inner product  $\langle \text{cat} | E^T E | \text{dog} \rangle \in \mathbb{R}^+$  to be strictly greater than zero: their contexts share at least the words “pet”, “food”, etc.

Then for every verb  $v \in R$ , we compute its meaning by summing over the outer products  $E | s \rangle \otimes E | o \rangle : 1 \rightarrow n \otimes n$  of all the encoded entities that it relates, i.e. all  $s, o \in E$  such that  $(s, v, o) \in K$ . Given an encoding matrix  $E$  and a knowledge graph  $K$  we store this construction as a matrix:



encoding every relation  $v$  in the *verb space*  $N \otimes N$ . Finally, we can compute the semantics of “subject verb object” sentences using the following recipe: let  $T = E \otimes R \otimes E$  be the encoding matrix for triples,

1. tensor the encoded states for  $t = (s, v, o) \in E \times R \times E$  together

$$T | t \rangle = E | s \rangle \otimes R | v \rangle \otimes E | o \rangle : 1 \rightarrow n \otimes (n \otimes n) \otimes n$$

2. translate the grammar into a linear map in  $\text{Mat}_{\mathbb{R}}^{\dagger}$ , here the effect

$$G = \text{cup}_n \otimes \text{cup}_n : (n \otimes n) \otimes (n \otimes n) \rightarrow 1$$


3. compose word meanings with grammar to get a scalar  $GT | t \rangle \in \mathbb{R}^+$ .

Computing these real-valued scalars allows to generalise to unseen sentences and outperform the standard bag-of-words approach on word disambiguation tasks, where the experimental data can be seen as a small knowledge graph [11]. We can give the semantics of more complex sentences — hence more complex knowledge graph patterns — by modeling relative pronouns as three-output spiders over  $n$ . For any  $x, y : 1 \rightarrow n$ , if we encode the word “that” as the Kronecker delta  $\delta_n^{0,3} : n^3 \rightarrow 1$  and apply the appropriate grammatical reduction we get:

$$G' ( x \otimes \text{that} \otimes y ) = \delta_n^{2,1} ( x \otimes y ) : 1 \rightarrow n$$


where  $G' = \text{cup}_n \otimes \text{id}_n \otimes \text{cup}_n : n^5 \rightarrow 1$ . This is the coordinate-wise multiplication of the vectors  $x$  and  $y$ , which can be seen as the intersection of entities in co-occurrence space, see [15] and [16].

**Example 1.** *If we encode some philosophy into our knowledge graph  $K$ , we want that:*



*i.e. all men are mortal.*

**Example 2.** *We compute the semantics of “Alice loves boys that tell jokes.” as the following scalar:*



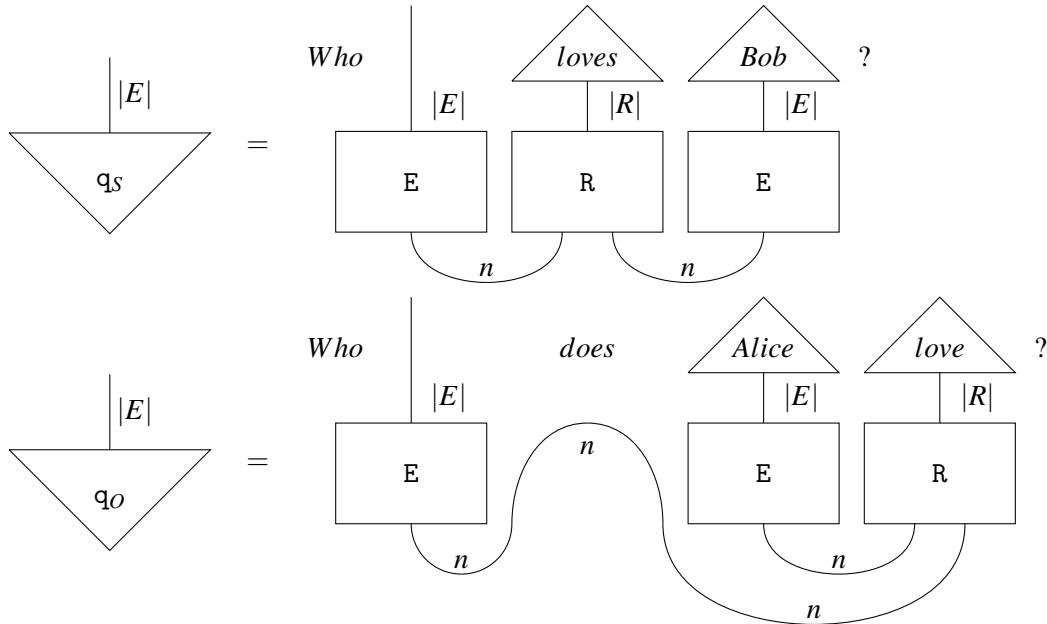
### 3 From Questions and Discourse to Knowledge Graph Queries

From an NLP perspective, computing the diagram for the sentence above can equivalently be seen as answering the question “Do the boys that Alice loves tell jokes?”. We will exploit this question-answer duality to show how natural language can be translated into machine-readable format in a structure-preserving way. Given a concrete implementation of a knowledge graph, also called a *triplestore*, the semantics of yes-no questions can be interpreted as the true - false output of ASK-type queries in the SPARQL protocol and RDF query language. As a first step towards more complex query patterns, we extend this language-to-database translation to object and subject questions as SELECT-type queries.


In [13], Lambek gives a thorough analysis of the grammar of questions in terms of pregroups. Here we will leave the subtleties of the English language aside and only assume designated grammatical types for object and subject questions; in our semantic category  $\text{Mat}(S)$ , both correspond to effects  $q : |E| \rightarrow 1$ . Indeed, given a potential answer  $a \in E$  the inner product  $q|a\rangle \in S$  gives us a measure of how well  $a$  answers the query  $q$ ; comparing these scalars allow us to translate a natural language question into a ranking of its possible answers. Note that if we interpret  $q$  with  $S = \mathbb{B}$  we get crisp answers while taking  $S = \mathbb{R}^+$  and the encoding  $T = E \otimes R \otimes E$  yields real-valued approximations.

In order to form a subject question, Lambek starts by constructing incomplete sentences with a hole in the subject position, e.g. “- loves Bob”. Object questions proceed dually, for example “Alice loves -”. Question words like “who” and “whom” can then be seen as processes from these incomplete sentences to


the question type. In our setup, the semantics of both words are given by the encoding matrix  $E : |E| \rightarrow n$ , they take the state for the answer as input and feed it in the computation of the corresponding query. The semantics for “Who loves Bob?” and “Who does Alice love?” are given by:



Note that the semantics of “does” is given by a cap, i.e. a maximally-entangled state of two noun systems, this is discussed further in [3]. Indeed, using the symmetry  $(\langle x| \otimes \langle y|) \text{cap}_{|E|} = (\langle y| \otimes \langle x|) \text{cap}_{|E|}$  for all  $x, y \in E$  and the snake equations we can rewrite object questions and see the cap serve as an information-passing mechanism:



We can extend this to two-variable SELECT queries such as



however, general SELECT queries can have an unbounded number of variables whereas natural language questions of three arguments or more will sound awkward to a human speaker — picture a stranger

asking you: “Who did what where with whom?”... Instead, we argue that natural language speakers can express complex query patterns through ambiguous *discourse* together with *anaphora* as a resource for interaction of meanings. Concretely, we translate personal pronouns such as “us” and “them” in terms of SPARQL query variables ranging over the set of entities  $E$ , then we say a sentence is *anaphoric* when it contains a personal pronoun: its meaning depends on that of another expression in context.

We define an *atomic sentence* as one of the form “subject verb object”, where both subject and object may be either entities or personal pronouns. In the anaphoric case we obtain a diagram with an open wire for the anaphoric expression, representing the input required from the context in order to compute the meaning of the sentence. Once the ambiguity has been resolved, i.e. once we have assigned an entity to the pronoun, we get a diagram with no open wires as in the unambiguous case of Section 2. Finally a *basic anaphoric discourse* is defined as a sequence of these atomic sentences, when it has  $k$  pronouns in total its semantics is an effect of dimension  $|E|^k$ : a knowledge graph query of  $k$  variables.

**Definition.** Take  $a, b \in \{0, 1\}$ , an atomic sentence is a matrix in

$$\left\{ G \left( \text{sub} \otimes R |v\rangle \otimes \text{obj} \right) : |E|^a \otimes |E|^b \rightarrow 1 \right\}$$

where the verb  $v$  ranges over the relations  $R$ . If  $a = 1$ , the sentence is anaphoric in the subject position and we take  $\text{sub} = E : |E| \rightarrow n$ . If  $a = 0$ ,  $\text{sub}$  ranges over the encoded entities  $E|s\rangle : 1 \rightarrow n$  and the sentence subject is unambiguous. Symmetrically for the object position:


$$\text{obj} = \begin{cases} E & : |E| \rightarrow n & \text{if } b = 1 \\ E|o\rangle & : 1 \rightarrow n & \text{for } o \in E \text{ otherwise} \end{cases}$$

We define a basic anaphoric discourse as the matrix  $d = \bigotimes_{i < l} \tau_i : |E|^k \rightarrow 1$  obtained by tensoring a list of  $l$  atomic sentences

$$\tau_i : |E|^{a_i} \otimes |E|^{b_i} \rightarrow 1$$

where  $k = \sum_{i < l} (a_i + b_i)$ , i.e. the discourse  $d$  contains  $k$  anaphoric expressions in total.

**Example 3.** The semantics of “Spinoza influenced him. He discovered calculus.” is given by the effect



## 4 Natural language understanding as a process

We have seen that the semantics of *basic anaphoric discourses* involving  $k$  anaphoric expressions can be given in terms of knowledge graph queries of  $k$  variables. However, interpreting directly the effect  $|E|^2 \rightarrow 1$  corresponding to the discourse in the example above as a SPARQL query of two variables will give counter-intuitive answers such as:

“Spinoza influenced Wittgenstein. Newton discovered calculus.”

whereas an English speaker will have naturally assumed that the two pronouns “him” and “he” have to refer to the same entity. Without this constraint, maximising over the complex query would be equivalent to maximising the outcome for the two sentences independently.

*Discourse representation theory* is a computational framework first introduced in [12] for manipulating such natural language constraints. In this paper, we model basic *discourse representation structures* (DRS) in the sense of [2], where entities are called “discourse referents”, atomic sentences “literals” and pronouns “variables”. Using the set of constraints formalised in terms of DRS, together with our encoding  $T = E \otimes R \otimes E$  allows us to formalise *probabilistic anaphora resolution* as an optimisation problem.

**Definition.** Given a matching function  $\mu : \{0, \dots, k-1\} \rightarrow E$  which assigns an entity in  $E$  to each pronoun in  $\mathfrak{d}$ , the resolution of the ambiguity in the discourse  $\mathfrak{d}$  is defined as the following scalar:

$$\mathcal{A}(\mathfrak{d}, \mu) = d \bigotimes_{i < k} | \mu(i) \rangle \in S$$

Given a set  $\mathcal{D}(\mathfrak{d}) \subseteq E^{\{0 \dots k-1\}}$  of matching functions satisfying some DRS constraints for the anaphoric discourse  $\mathfrak{d}$ , we define probabilistic anaphora resolution as the following optimisation problem:

$$\operatorname{argmax}_{\mu \in \mathcal{D}(\mathfrak{d})} \mathcal{A}(\mathfrak{d}, \mu)$$

The function  $\mu$  can only give a static picture of the ambiguity resolution’s result, instead we want a fine-grained diagram of the computational process involved. We give a resource-aware reformulation of anaphora resolution in terms of an *entity-store*, defined as the state:

$$!E = \bigotimes_{e \in E} |e\rangle : 1 \rightarrow |E|^{|E|}$$

From the matching function  $\mu$  we construct a *matching process*  $!\mu$  which takes an entity-store  $!E$  and feeds the appropriate entities into the ambiguous sentences in  $\mathfrak{d}$ . Graphically this construction amounts to drawing the graph of the function  $\mu$ , labeling all the wires with  $|E|$  and interpreting this as a diagram in  $\operatorname{Mat}(S)$ . When  $i$  edges of the graph meet at a vertex in  $E$  — i.e. when the same entity is referenced by  $i$  distinct pronouns — we witness this with a Kronecker delta: the head of a 1-input  $i$ -output spider. Whereas in section 2 the multi-input one-output spider acted as intersection of meanings, here the dual process copies an entity state and feeds it to its  $i$  outputs i.e. for all  $e \in E$  we have that:

$$\begin{aligned} (\delta_{|E|}^{1,i}) |e\rangle &= \sum_{e' \in E} |e'\rangle^{(\otimes i)} \langle e'|e\rangle \\ &= \left( \sum_{e' \neq e} |e'\rangle^{(\otimes i)} \underbrace{\langle e'|e\rangle}_{=0} \right) + |e\rangle^{(\otimes i)} \underbrace{\langle e|e\rangle}_{=1} = |e\rangle^{(\otimes i)} \end{aligned}$$

**Theorem 2.** Given a matching function  $\mu : \{0, \dots, k-1\} \rightarrow E$  we can construct a matching process  $!\mu : |E|^{|E|} \rightarrow |E|^k$  such that:

$$\forall \mathfrak{d} : |E|^k \rightarrow 1 \cdot \mathcal{A}(\mathfrak{d}, \mu) = \mathfrak{d}(!\mu)(!E)$$

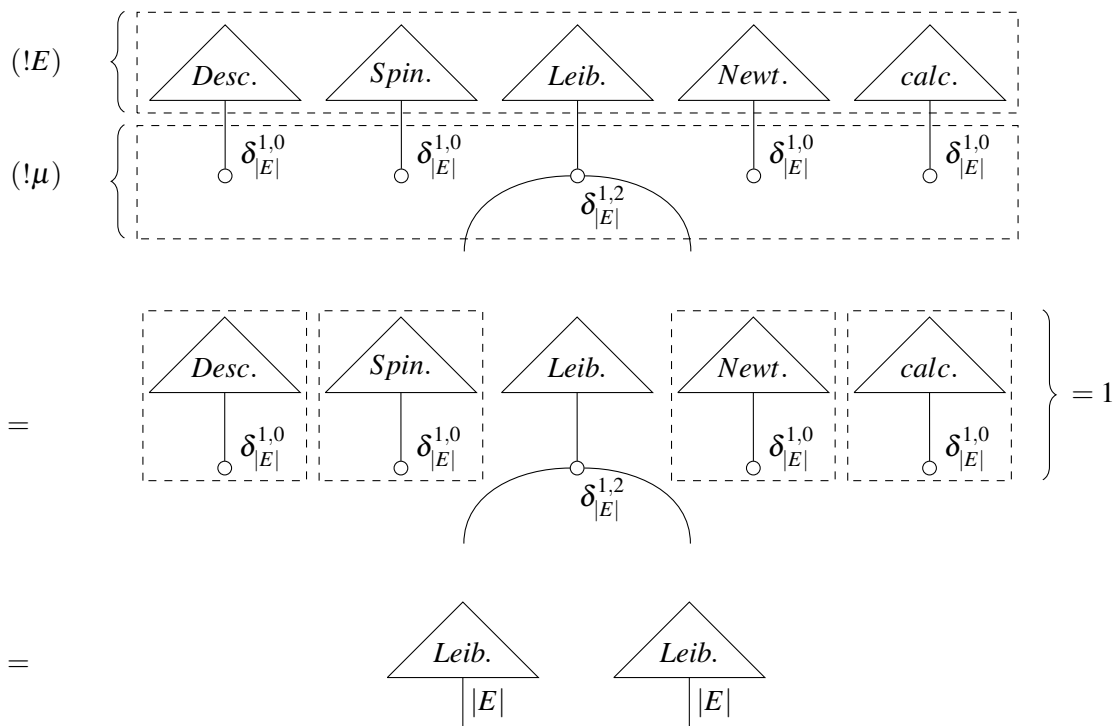
*Proof.* In set-theoretic terms,  $!E$  is the list of all the elements  $e \in E$ , ordered by the indices of the corresponding vectors  $|e\rangle : 1 \rightarrow |E|$ . Then  $!\mu$  is the incidence matrix of the graph for the function

$$f_\mu ( e_0, \dots, e_{|E|-1} ) = ( e_{\mu(0)}, \dots, e_{\mu(k-1)} ) : E^{|E|} \rightarrow E^k$$



In the case when the graph of the function  $f_\mu$  is not planar, constructing  $!\mu$  as a diagram in  $\text{Mat}(S)$  requires a formalisation of the intuitive notion of “wire-crossing”, which can be obtained using the *symmetry morphisms* of  $\text{Mat}(S)$ . This falls outside the scope of our discussion but we refer the interested reader to any introduction to *symmetric monoidal categories* (e.g. [9]) which are the suitable framework in this context.  $\square$


**Example 4.** Take  $E = \{\text{Descartes}, \text{Spinoza}, \text{Leibniz}, \text{Newton}, \text{calculus}\}$ , let  $\mu$  be the constant function  $\{0, 1 \mapsto \text{Leibniz}\}$ . We have:



Even though our construction involves systems of exponential dimensions, the concrete computation it induces is linear in the size  $n \times |E|$  of the encoding matrix: we are only retrieving and copying vectors. In effect, what this abstract formulation allows is to manipulate the knowledge graph, the entity encoding and matching explicitly, while hiding the bureaucratic process of labeling, indexing and wiring the entity-vectors required to compute complex queries.

We can now focus on what really is the core of the model: the interaction between the symbolic world of RDF and the statistical world of distributional semantics. In diagrammatic terms, this is witnessed by the interaction of distinct spiders, i.e. Kronecker deltas over two different dimensions:  $\delta_{|E|}$  for copying and matching entities,  $\delta_n$  for their intersection in the noun space  $N$ .

**Example 5.** We compute the resolution  $\mathcal{A}(d, \mu) = d(!\mu)(!E) \in \mathbb{R}$  as the scalar:



If we extract our encoding matrix  $E$  and knowledge graph  $K$  from an encyclopedia, we should expect this scalar to be close to 1. Choosing  $\mu = \{0, 1 \mapsto \text{Descartes}\}$  instead should make it close to 0.

In order to get the class of all basic DRS, we need to assume that for every  $r \in R$  we have its negation  $\text{not}(r) \in R$  such that  $E(x, \text{not}(r), y) \in K$  if and only if  $(x, r, y) \notin K$ . Lifting this function  $\text{not} : R \rightarrow R$  to the matrix of its graph  $|\text{not}\rangle : |R| \rightarrow |R|$  and deriving a compositional semantics for natural language negations in  $\text{Mat}_S$  from this is left for future work. Nevertheless, if we restrict ourselves to the negation-free fragment, we can already give a semantics for basic anaphoric discourses in terms of the conjunctive fragment of SPARQL. Indeed, we can generate any query of the form

$$\text{SELECT } * \text{ WHERE } \{ \langle x, r, y \rangle \in \text{BGP} \}$$

where  $x$  and  $y$  are taken to range over both entities and variables. In the Semantic Web community, the set of triples BGP is also called a *basic graph pattern* [19].

## Conclusion

We have developed a process-theoretic framework for natural language understanding, making use of both statistical and logical resources to translate question answering and anaphora resolution as probabilistic knowledge base queries. Using formal diagrams of matrices we give a high-level picture of NLP algorithms which are motivated both from the point of view of AI and of human cognition.

Even though we focused on  $S = \mathbb{B}$  and  $S = \mathbb{R}^+$ , our discussion can be extended to any other semi-ring structure such as the unit interval  $[0, 1]$  with  $\text{min}$  and  $\text{max}$  encoding fuzzy truth values. It is also possible to incorporate further structure such as the convexity property used in conceptual space models of cognition, or taking proof relevance into account and record *why* are entities related [7].

In the RDF language, variables can also serve as *blank nodes* also called *anonymous resources*: entities that are used in the computation but do not appear in the output of the query, e.g.  $y$  in

$$\text{SELECT } x \text{ WHERE } \{ \langle x, \text{loves}, y \rangle, \langle y, \text{loves}, x \rangle \}$$

In  $\text{Mat}_S$ , we conjecture that these would translate in terms of an *entity-discarding* process  $\delta_{|E|}^{0,1} : 1 \rightarrow |E|$ .

The translation from pregroup grammars to RDF triples is developed further in [10], in order to model count nouns (e.g. “a cat”) and relative pronouns (e.g. “whose tail”), Delpeuch and Preller introduce a notion of categories with *side effects* — i.e. creation and update of blank nodes. It would be interesting to investigate the connections with our current work.

We would like to investigate how our model relates to concrete architectures for NLP, such as the Tensor Product Recurrent Networks (TPRNs) of [14]. Another machine learning model which would fit our approach is that of [20]: this method uses gradient descent over a landscape of complex-valued encoding matrices  $E$  and  $R$  to predict missing values in a knowledge graph  $K$ . Linguistically, the passage to the alternative semi-ring of complex numbers would allow to give two distinct encodings for entities: in our model the object-pronoun “them” would be modelled as the complex conjugate of the subject “they”. Modelling the active-to-passive switch with diagrams of complex matrices, as well as exploring potential implementations making use of quantum resources is left to later work.

In this paper we looked only at one discourse  $d$  at a time, assuming a global matching process  $!μ$  and entity-store  $!E$ . Using methods from sheaf theory would allow to study the passage from local to global ambiguity and investigate the notion of *contextuality* common to natural language, quantum mechanics and database theory [1]. This finer-grained analysis of the resolution process, as well as the implementation of other DRS operations, such as disjunction, implication and quantification, in the category  $\text{Mat}_{\mathcal{S}}$  is left for future work.

## References

- [1] Samson Abramsky (2016): *Contextual semantics: From quantum mechanics to logic, databases, constraints, and complexity*. In: *Contextuality from Quantum Physics to Psychology*, World Scientific, pp. 23–50, DOI: 10.1142/9789814730617\_0002.
- [2] Samson Abramsky & Mehrnoosh Sadrzadeh (2014): *Semantic Unification*. In: *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 1–13, DOI: 10.1007/978-3-642-54789-8\_1.
- [3] Stephen Clark, Bob Coecke, Edward Grefenstette, Stephen Pulman & Mehrnoosh Sadrzadeh (2014): *A quantum teleportation inspired algorithm produces sentence meaning from word meaning and grammatical structure*. *Malaysian Journal of Mathematical Sciences* 8, pp. 15–25. ArXiv:1305.0556.
- [4] Stephen Clark, Bob Coecke & Mehrnoosh Sadrzadeh (2008): *A Compositional Distributional Model of Meaning*. *Proceedings of the Second Symposium on Quantum Interaction (QI-2008)*, pp. 133–140.
- [5] Stephen Clark, Bob Coecke & Mehrnoosh Sadrzadeh (2010): *Mathematical foundations for a compositional distributional model of meaning*. In J. van Benthem, M. Moortgat & W. Buszkowski, editors: *A Festschrift for Jim Lambek*, *Linguistic Analysis* 36, pp. 345–384. Arxiv:1003.4394.
- [6] Bob Coecke & Ross Duncan (2011): *Interacting quantum observables: categorical algebra and diagrammatics*. *New Journal of Physics* 13, p. 043016, DOI: 10.1088/1367-2630/13/4/043016. arXiv:quant-ph/09064725.
- [7] Bob Coecke, Fabrizio Genovese, Martha Lewis, Dan Marsden & Alex Toumi (2018): *Generalized relations in linguistics & cognition*. *Theoretical Computer Science*, DOI: 10.1016/j.tcs.2018.03.008. Available at <http://www.sciencedirect.com/science/article/pii/S0304397518301476>.
- [8] Bob Coecke & Aleks Kissinger (2017): *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press. Available at <http://dx.doi.org/10.1017/9781316219317>.

- [9] Bob Coecke & Eric Oliver Paquette (2011): *Categories for the practicing physicist*. In B. Coecke, editor: *New Structures for Physics*, Lecture Notes in Physics, Springer, pp. 167–271, DOI: 10.1007/978-3-642-12821-9. arXiv:0905.3010.
- [10] Antonin Delpeuch & Anne Preller (2014): *From Natural Language to RDF Graphs with Pregroups*. In: *EACL'2014: 14th Conference of the European Chapter of the Association for Computational Linguistics*, EACL, Gothenburg, Sweden, pp. 55–62. Available at <https://hal-lirmm.ccsd.cnrs.fr/lirmm-00992381>.
- [11] Edward Grefenstette & Mehrnoosh Sadrzadeh (2011): *Experimental Support for a Categorical Compositional Distributional Model of Meaning*. In: *The 2014 Conference on Empirical Methods on Natural Language Processing*, pp. 1394–1404. ArXiv:1106.4058.
- [12] Hans Kamp & Uwe Reyle (1993): *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. *Studies in Linguistics and Philosophy* 42, Springer, Dordrecht, DOI: 10.1007/978-94-017-1616-1.
- [13] Jim Lambek (2008): *From Word to Sentence: A Computational Algebraic Approach to Grammar*. Open access publications, Polimetrica.
- [14] Hamid Palangi, Paul Smolensky, Xiaodong He & Li Deng (2017): *Question-Answering with Grammatically-Interpretable Representations*. arXiv preprint arXiv:1705.08432.
- [15] Mehrnoosh Sadrzadeh, Stephen Clark & Bob Coecke (2013): *The Frobenius anatomy of word meanings I: subject and object relative pronouns*. *Journal of Logic and Computation* 23, pp. 1293–1317, DOI: 10.1093/logcom/ext044. ArXiv:1404.5278.
- [16] Mehrnoosh Sadrzadeh, Stephen Clark & Bob Coecke (2014): *The Frobenius anatomy of word meanings II: possessive relative pronouns*. *Journal of Logic and Computation* abs/1406.4690, p. exu027, DOI: 10.1093/logcom/exu027. Available at <http://arxiv.org/abs/1406.4690>.
- [17] P. Selinger (2010): *A Survey of Graphical Languages for Monoidal Categories*. In: *New Structures for Physics*, Springer Berlin Heidelberg, pp. 289–355, DOI: 10.1007/978-3-642-12821-9\_4.
- [18] Paul Smolensky (1987): *Connectionist AI, symbolic AI, and the brain*. *Artif. Intell. Rev.* 1(2), pp. 95–109, DOI: 10.1007/BF00130011.
- [19] Giorgio Stefanoni, Boris Motik & Egor V. Kostylev (2018): *Estimating the Cardinality of Conjunctive Queries over RDF Data Using Graph Summarisation*. In Pierre-Antoine Champin, Fabien L. Gandon, Mounia Lalmas & Panagiotis G. Ipeirotis, editors: *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, ACM, pp. 1043–1052, DOI: 10.1145/3178876.3186003.
- [20] Théo Trouillon, Christopher R. Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel & Guillaume Bouchard (2017): *Knowledge Graph Completion via Complex Tensor Factorization*. *J. Mach. Learn. Res.* 18(1), pp. 4735–4772. Available at <http://dl.acm.org/citation.cfm?id=3122009.3208011>.