

The Frobenius Anatomy of Relative Pronouns

Stephen Clark
University of Cambridge
Computer Laboratory
sc609@cam.ac.uk

Bob Coecke
University of Oxford
Dept. of Computer Science
coecke@cs.ox.ac.uk

Mehrnoosh Sadrzadeh
Queen Mary, University of London
School of Electronic
Engineering and Computer Science
mehrnoosh.sadrzadeh@eecs.qmul.ac.uk

Abstract

This paper develops a compositional vector-based semantics of relative pronouns within a categorical framework. Frobenius algebras are used to formalise the operations required to model the semantics of relative pronouns, including passing information between the relative clause and the modified noun phrase, as well as copying, combining, and discarding parts of the relative clause. We develop two instantiations of the abstract semantics, one based on a truth-theoretic approach and one based on corpus statistics.

1 Introduction

Ordered algebraic structures and sequent calculi have been used extensively in Computer Science and Mathematical Logic. They have also been used to formalise and reason about natural language. Lambek (1958) used the ordered algebra of residuated monoids to model grammatical types, their juxtapositions and reductions. Relational words such as verbs have implicative types and are modelled using the residuals to the monoid multiplication. Later, Lambek (1999) simplified these algebras in favour of *pregroups*. Here, there are no binary residual operations, but each element of the algebra has a left and a right residual.

In terms of semantics, pregroups do not naturally lend themselves to a model-theoretic treatment (Montague, 1974). However, pregroups are suited to a radically different treatment of semantics, namely distributional semantics (Schütze, 1998). Distributional semantics uses vector spaces based on contextual co-occurrences to model the meanings of words. Coecke et al. (2010) show how a *compositional* semantics can be developed within a vector-based framework, by exploiting the fact that vector spaces with linear maps and

pregroups both have a compact closed categorical structure (Kelly and Laplaza, 1980; Preller and Lambek, 2007). Some initial attempts at implementation include Grefenstette and Sadrzadeh (2011a) and Grefenstette and Sadrzadeh (2011b).

One problem with the distributional approach is that it is difficult to see how the meanings of some words — e.g. logical words such as *and*, *or*, and relative pronouns such as *who*, *which*, *that*, *whose* — can be modelled contextually. Our focus in this paper is on relative pronouns in the distributional compositional setting.

The difficulty with pronouns is that the contexts in which they occur do not seem to provide a suitable representation of their meanings: pronouns tend to occur with a great many nouns and verbs. Hence, if one applies the contextual co-occurrence methods of distributional semantics to them, the result will be a set of dense vectors which do not discriminate between different meanings. The current state-of-the-art in compositional distributional semantics either adopts a simple method to obtain a vector for a sequence of words, such as adding or multiplying the contextual vectors of the words (Mitchell and Lapata, 2008), or, based on the grammatical structure, builds linear maps for some words and applies these to the vector representations of the other words in the string (Baroni and Zamparelli, 2010; Grefenstette and Sadrzadeh, 2011a). Neither of these approaches produce vectors which provide a good representation for the meanings of relative clauses.

In the grammar-based approach, one has to assign a linear map to the relative pronoun, for instance a map f as follows:

$$\overrightarrow{\text{men who like Mary}} = f(\overrightarrow{\text{men}}, \overrightarrow{\text{like Mary}})$$

However, it is not clear what this map should be. Ideally, we do not want it to depend on the frequency of the co-occurrence of the relative pronoun with the relevant basis vectors. But both

of the above mentioned approaches rely heavily on the information provided by a corpus to build their linear maps. The work of Baroni and Zamparelli (2010) uses linear regression and approximates the context vectors of phrases in which the target word has occurred, and the work of Grefenstette and Sadrzadeh (2011a) uses the sum of Kronecker products of the arguments of the target word across the corpus.

The semantics we develop for relative pronouns and clauses uses the general operations of a Frobenius algebra over vector spaces (Coecke et al., 2008) and the structural categorical morphisms of vector spaces. We do not rely on the co-occurrence frequencies of the pronouns in a corpus and only take into account the structural roles of the pronouns in the meaning of the clauses. The computations of the algebra and vector spaces are depicted using string diagrams (Joyal and Street, 1991), which depict the interactions that occur among the words of a sentence. In the particular case of relative clauses, they visualise the role of the relative pronoun in passing information between the clause and the modified noun phrase, as well as copying, combining, and even discarding parts of the relative clause.

We develop two instantiations of the abstract semantics, one based on a truth-theoretic approach, and one based on corpus statistics, where for the latter the categorical operations are instantiated as matrix multiplication and vector component-wise multiplication. As a result, we will obtain the following for the meaning of a subject relative clause:

$$\overrightarrow{\text{men who like Mary}} = \overrightarrow{\text{men}} \odot (\overrightarrow{\text{love}} \times \overrightarrow{\text{Mary}})$$

The rest of the paper introduces the categorical framework, including the formal definitions relevant to the use of Frobenius algebras, and then shows how these structures can be used to model relative pronouns within the compositional vector-based setting.

2 Compact Closed Categories and Frobenius Algebras

This section briefly reviews compact closed categories and Frobenius algebras. For a formal presentation, see (Kelly and Laplaza, 1980; Kock, 2003; Baez and Dolan, 1995), and for an informal introduction see Coecke and Paquette (2008).

A compact closed category has objects A, B ; morphisms $f: A \rightarrow B$; a monoidal tensor $A \otimes B$

that has a unit I ; and for each object A two objects A^r and A^l together with the following morphisms:

$$\begin{aligned} A \otimes A^r &\xrightarrow{\epsilon_A^r} I \xrightarrow{\eta_A^r} A^r \otimes A \\ A^l \otimes A &\xrightarrow{\epsilon_A^l} I \xrightarrow{\eta_A^l} A \otimes A^l \end{aligned}$$

These morphisms satisfy the following equalities, sometimes referred to as the *yanking* equalities, where 1_A is the identity morphism on object A :

$$\begin{aligned} (1_A \otimes \epsilon_A^l) \circ (\eta_A^l \otimes 1_A) &= 1_A \\ (\epsilon_A^r \otimes 1_A) \circ (1_A \otimes \eta_A^r) &= 1_A \\ (\epsilon_A^l \otimes 1_A) \circ (1_{A^l} \otimes \eta_A^l) &= 1_{A^l} \\ (1_{A^r} \otimes \epsilon_A^r) \circ (\eta_A^r \otimes 1_{A^r}) &= 1_{A^r} \end{aligned}$$

A pregroup is a partial order compact closed category, which we refer to as *Preg*. This means that the objects of *Preg* are elements of a partially ordered monoid, and between any two objects $p, q \in \text{Preg}$ there exists a morphism of type $p \rightarrow q$ iff $p \leq q$. Compositions of morphisms are obtained by transitivity and the identities by reflexivity of the partial order. The tensor of the category is the monoid multiplication, and the epsilon and eta maps are as follows:

$$\begin{aligned} \epsilon_p^r &= p \cdot p^r \leq 1 & \eta_p^r &= 1 \leq p^r \cdot p \\ \epsilon_p^l &= p^l \cdot p \leq 1 & \eta_p^l &= 1 \leq p \cdot p^l \end{aligned}$$

Finite dimensional vector spaces and linear maps also form a compact closed category, which we refer to as *FVect*. Finite dimensional vector spaces V, W are objects of this category; linear maps $f: V \rightarrow W$ are its morphisms with composition being the composition of linear maps. The tensor product $V \otimes W$ is the linear algebraic tensor product, whose unit is the scalar field of vector spaces; in our case this is the field of reals \mathbb{R} . As opposed to the tensor product in *Preg*, the tensor between vector spaces is symmetric; hence we have a natural isomorphism $V \otimes W \cong W \otimes V$. As a result of the symmetry of the tensor, the two adjoints reduce to one and we obtain the following isomorphism:

$$V^l \cong V^r \cong V^*$$

where V^* is the dual of V . When the basis vectors of the vector spaces are fixed, it is further the case that the following isomorphism holds as well:

$$V^* \cong V$$

Elements of vector spaces, i.e. vectors, are represented by morphisms from the unit of tensor to their corresponding vector space; that is $\vec{v} \in V$ is represented by the morphism $\mathbb{R} \xrightarrow{\vec{v}} V$; by linearity this morphism is uniquely defined when setting $1 \mapsto \vec{v}$.

Given a basis $\{r_i\}_i$ for a vector space V , the epsilon maps are given by the inner product extended by linearity; i.e. we have:

$$\epsilon^l = \epsilon^r : V^* \otimes V \rightarrow \mathbb{R}$$

given by:

$$\sum_{ij} c_{ij} \psi_i \otimes \phi_j \mapsto \sum_{ij} c_{ij} \langle \psi_i | \phi_j \rangle$$

Similarly, eta maps are defined as follows:

$$\eta^l = \eta^r : \mathbb{R} \rightarrow V \otimes V^*$$

and are given by:

$$1 \mapsto \sum_i r_i \otimes r_i$$

A Frobenius algebra in a monoidal category $(\mathcal{C}, \otimes, I)$ is a tuple $(X, \Delta, \iota, \mu, \zeta)$ where, for X an object of \mathcal{C} , the triple (X, Δ, ι) is an internal comonoid; i.e. the following are coassociative and counital morphisms of \mathcal{C} :

$$\Delta : X \rightarrow X \otimes X \quad \iota : X \rightarrow I$$

Moreover (X, μ, ζ) is an internal monoid; i.e. the following are associative and unital morphisms:

$$\mu : X \otimes X \rightarrow X \quad \zeta : I \rightarrow X$$

And finally the Δ and μ morphisms satisfy the following *Frobenius condition*:

$$(\mu \otimes 1_X) \circ (1_X \otimes \Delta) = \Delta \circ \mu = (1_X \otimes \mu) \circ (\Delta \otimes 1_X)$$

Informally, the comultiplication Δ decomposes the information contained in one object into two objects, and the multiplication μ combines the information of two objects into one.

Frobenius algebras were originally introduced in the context of representation theorems for group theory (Frobenius, 1903). Since then, they have found applications in other fields of mathematics and physics, e.g. in topological quantum field theory (Kock, 2003). The above general categorical definition is due to Carboni and Walters (1987). In

what follows, we use Frobenius algebras that characterise vector space bases (Coecke et al., 2008).

In the category of finite dimensional vector spaces and linear maps $FVect$, any vector space V with a fixed basis $\{\vec{v}_i\}_i$ has a Frobenius algebra over it, explicitly given by:

$$\begin{aligned} \Delta &:: \vec{v}_i \mapsto \vec{v}_i \otimes \vec{v}_i & \iota &:: \vec{v}_i \mapsto 1 \\ \mu &:: \vec{v}_i \otimes \vec{v}_j \mapsto \delta_{ij} \vec{v}_i & \zeta &:: 1 \mapsto \sum_i \vec{v}_i \end{aligned}$$

where δ_{ij} is the Kronecker delta.

Frobenius algebras over vector spaces with orthonormal bases are moreover *isometric* and *commutative*. A commutative Frobenius Algebra satisfies the following two conditions for $\sigma : X \otimes Y \rightarrow Y \otimes X$, the symmetry morphism of $(\mathcal{C}, \otimes, I)$:

$$\sigma \circ \Delta = \Delta \quad \mu \circ \sigma = \mu$$

An isometric Frobenius Algebra is one that satisfies the following axiom:

$$\mu \circ \Delta = 1$$

The vector spaces of distributional models have fixed orthonormal bases; hence they have isometric commutative Frobenius algebras over them.

The comultiplication Δ of an isometric commutative Frobenius Algebra over a vector space encodes vectors of lower dimensions into vectors of higher dimensional tensor spaces; this operation is referred to as *copying*. In linear algebraic terms, $\Delta(\vec{v}) \in V \otimes V$ is a diagonal matrix whose diagonal elements are weights of $\vec{v} \in V$. The corresponding multiplication μ encodes vectors of higher dimensional tensor spaces into lower dimensional spaces; this operation is referred to as *combining*. For $\vec{w} \in V \otimes V$, we have that $\mu(\vec{w}) \in V$ is a vector consisting only of the diagonal elements of \vec{w} .

As a concrete example, take V to be a two dimensional space with basis $\{\vec{v}_1, \vec{v}_2\}$; then the basis of $V \otimes V$ is $\{\vec{v}_1 \otimes \vec{v}_1, \vec{v}_1 \otimes \vec{v}_2, \vec{v}_2 \otimes \vec{v}_1, \vec{v}_2 \otimes \vec{v}_2\}$. For a vector $v = a\vec{v}_1 + b\vec{v}_2$ in V we have:


$$\Delta(v) = \Delta \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} = a \vec{v}_1 \otimes \vec{v}_1 + b \vec{v}_2 \otimes \vec{v}_2$$

And for a matrix $w = a \vec{v}_1 \otimes \vec{v}_1 + b \vec{v}_1 \otimes \vec{v}_2 + c \vec{v}_2 \otimes \vec{v}_1 + d \vec{v}_2 \otimes \vec{v}_2$ in $V \otimes V$, we have:


$$\mu(w) = \mu \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} a \\ d \end{pmatrix} = a \vec{v}_1 + d \vec{v}_2$$

3 String Diagrams


The framework of compact closed categories and Frobenius algebras comes with a complete diagrammatic calculus that visualises derivations, and which also simplifies the categorical and vector space computations. Morphisms are depicted by boxes and objects by lines, representing their identity morphisms. For instance a morphism $f: A \rightarrow B$, and an object A with the identity arrow $1_A: A \rightarrow A$, are depicted as follows:




The tensor products of the objects and morphisms are depicted by juxtaposing their diagrams side by side, whereas compositions of morphisms are depicted by putting one on top of the other; for instance the object $A \otimes B$, and the morphisms $f \otimes g$ and $f \circ h$, for $f: A \rightarrow B, g: C \rightarrow D$, and $h: B \rightarrow C$, are depicted as follows:




The ϵ maps are depicted by cups, η maps by caps, and yanking by their composition and straightening of the strings. For instance, the diagrams for $\epsilon^l: A^l \otimes A \rightarrow I, \eta: I \rightarrow A \otimes A^l$ and $(\epsilon^l \otimes 1_A) \circ (1_A \otimes \eta^l) = 1_A$ are as follows:




The composition of the ϵ and η maps with other morphisms is depicted as before, that is by juxtaposing them one above the other. For instance the diagrams for the compositions $(1_{B^l} \otimes f) \circ \epsilon^l$ and $\eta^l \circ (1_{A^l} \otimes f)$ are as follows:




As for Frobenius algebras, the diagrams for the monoid and comonoid morphisms are as follows:



with the Frobenius condition being depicted as:




The defining axioms guarantee that any picture depicting a Frobenius computation can be reduced to a normal form that only depends on the number of input and output strings of the nodes, independent of the topology. These normal forms can be simplified to so-called ‘spiders’:




In the category $FVect$, apart from spaces V, W , which are objects of the category, we also have vectors \vec{v}, \vec{w} . These are depicted by their representing morphisms and as triangles with a number of strings emanating from them. The number of strings of a triangle denote the tensor rank of the vector; for instance, the diagrams for $\vec{v} \in V, \vec{v}^l \in V \otimes W$, and $\vec{v}^l \in V \otimes W \otimes Z$ are as follows:



Application of a linear map to a vector is depicted using composition of their corresponding morphisms. For instance, for $f: V \rightarrow W$ and $\vec{v} \in V$, the application $f(\vec{v})$ is depicted by the composition $I \xrightarrow{\vec{v}} V \xrightarrow{f} W$.



Applications of the Frobenius maps to vectors are depicted in a similar fashion; for instance $\mu(\vec{v} \otimes \vec{v})$ is the composition $I \otimes I \xrightarrow{\vec{v} \otimes \vec{v}} V \otimes V \xrightarrow{\mu} V$ and $\iota(\vec{v})$ is the composition $I \xrightarrow{\vec{v}} V \xrightarrow{\iota} I$, depicted as follows:




4 Vector Space Interpretations

The grammatical structure of a language is encoded in the category *Preg*: objects are grammatical types (assigned to words of the language) and morphisms are grammatical reductions (encoding the grammatical formation rules of the language). For instance, the grammatical structure of the sentence “Men love Mary” is encoded in the assignment of types n to the noun phrases “men” and “Mary” and $n^r \otimes s \otimes n^l$ to the verb “love”, and in the reduction map $\epsilon_n^l \otimes 1_s \otimes \epsilon_n^r$. The application of this reduction map to the tensor product of the word types in the sentence results in the type s :

$$(\epsilon_n^l \otimes 1_s \otimes \epsilon_n^r)(n \otimes (n^r \otimes s \otimes n^l) \otimes n) \rightarrow s$$

To each reduction map corresponds a string diagram that depicts the structure of reduction:



In Coecke et al. (2010) the pregroup types and reductions are interpreted as vector spaces and linear maps, achieved via a homomorphic mapping from *Preg* to *FVect*. Categorically speaking, this map is a strongly monoidal functor:

$$F: \text{Preg} \rightarrow \text{FVect}$$

It assigns vector spaces to the basic types as follows:

$$F(1) = I \quad F(n) = N \quad F(s) = S$$

and to the compound types by monoidality as follows; for x, y objects of *Preg*:

$$F(x \otimes y) = F(x) \otimes F(y)$$

Monoidal functors preserve the compact structure; that is the following holds:

$$F(x^l) = F(x^r) = F(x)^*$$

For instance, the interpretation of a transitive verb is computed as follows:

$$\begin{aligned} F(n^r \otimes s \otimes n^l) &= F(n^r) \otimes F(s) \otimes F(n^l) = \\ &F(n)^* \otimes F(s) \otimes F(n)^* = N \otimes S \otimes N \end{aligned}$$

This interpretation means that the meaning vector of a transitive verb is a vector in $N \otimes S \otimes N$.


The pregroup reductions, i.e. the partial order morphisms of *Preg*, are interpreted as linear maps: whenever $p \leq q$ in *Preg*, we have a linear map $f_{\leq}: F(p) \rightarrow F(q)$. The ϵ and η maps of *Preg* are interpreted as the ϵ and η maps of *FVect*. For instance, the pregroup reduction of a transitive verb sentence is computed as follows:

$$\begin{aligned} F(\epsilon_n^r \otimes 1_s \otimes \epsilon_n^l) &= F(\epsilon_n^r) \otimes F(1_s) \otimes F(\epsilon_n^l) = \\ &F(\epsilon_n)^* \otimes F(1_s) \otimes F(\epsilon_n)^* = \epsilon_N \otimes 1_S \otimes \epsilon_N \end{aligned}$$

The distributional meaning of a sentence is obtained by applying the interpretation of the pregroup reduction of the sentence to the tensor product of the distributional meanings of the words in the sentence. For instance, the distributional meaning of “Men love Mary” is as follows:

$$F(\epsilon_n^r \otimes 1_s \otimes \epsilon_n^l)(\overrightarrow{\text{Men}} \otimes \overrightarrow{\text{love}} \otimes \overrightarrow{\text{Mary}})$$

This meaning is depictable via the following string diagram:



The next section applies these techniques to the distributional interpretation of pronouns. The interpretations are defined using: ϵ maps, for application of the semantics of one word to another; η maps, to pass information around by bridging intermediate words; and Frobenius operations, for copying and combining the noun vectors and discarding the sentence vectors.

5 Modelling Relative Pronouns

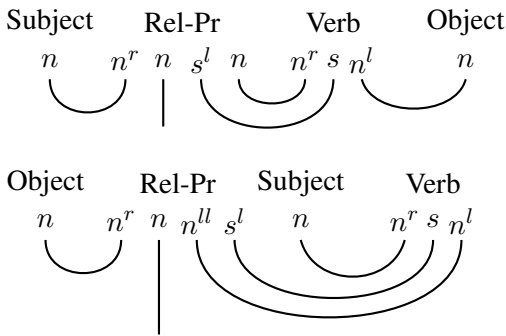
In this paper we focus on the subject and object relative pronouns, *who(m)*, *which* and *that*. Examples of noun phrases with subject relative pronouns are “men who love Mary”, “dog which ate cats”. Examples of noun phrases with object relative pronouns are “men whom Mary loves”, “book

that John read”. In the final example, “book” is the head noun, modified by the relative clause “that John read”. The intuition behind the use of Frobenius algebras to model such cases is the following. In “book that John read”, the relative clause acts on the noun (modifies it) via the relative pronoun, which passes information from the clause to the noun. The relative clause is then discarded, and the modified noun is returned. Frobenius algebras provide the machinery for all of these operations.

The pregroup types of the relative pronouns are as follows:

$$\begin{aligned} n^r n s^l n & \text{ (subject)} \\ n^r n n^l s^l & \text{ (object)} \end{aligned}$$


These types result in the following reductions:



The meaning spaces of these pronouns are computed using the mechanism described above:

$$\begin{aligned} F(n^r n s^l n) &= F(n^r) \otimes F(n) \otimes F(s^l) \otimes F(n) \\ &= N \otimes N \otimes S \otimes N \\ F(n^r n n^l s^l) &= F(n^r) \otimes F(n) \otimes F(n^l) \otimes F(s^l) \\ &= N \otimes N \otimes N \otimes S \end{aligned}$$


The semantic roles that these pronouns play are reflected in their categorical vector space meanings, depicted as follows:




with the following corresponding morphisms:

$$\begin{aligned} \text{Subj: } & (1_N \otimes \mu_N \otimes \zeta_S \otimes 1_N) \circ (\eta_N \otimes \eta_N) \\ \text{Obj: } & (1_N \otimes \mu_N \otimes 1_N \otimes \zeta_S) \circ (\eta_N \otimes \eta_N) \end{aligned}$$

The diagram of the meaning vector of the subject relative clause interacting with the head noun is as follows:

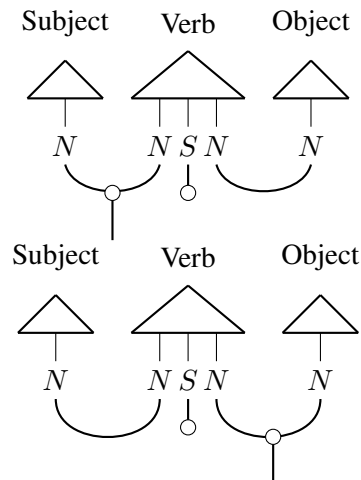


The diagram for the object relative clause is:



These diagrams depict the flow of information in a relative clause and the semantic role of its relative pronoun, which 1) passes information from the clause to the head noun via the η maps; 2) acts on the noun via the μ map; 3) discards the clause via the ζ map; and 4) returns the modified noun via 1_N . The ϵ maps pass the information of the subject and object nouns to the verb and to the relative pronoun to be acted on. Note that there are two different flows of information in these clauses: the ones that come from the grammatical structure and are depicted by ϵ maps (at the bottom of the diagrams), and the ones that come from the semantic role of the pronoun and are depicted by η maps (at the top of the diagrams).

The normal forms of these diagrams are:



Symbolically, they correspond to the following morphisms:

$$\begin{aligned} & (\mu_N \otimes \iota_S \otimes \epsilon_N) \left(\overrightarrow{\text{Subject}} \otimes \overrightarrow{\text{Verb}} \otimes \overrightarrow{\text{Object}} \right) \\ & (\epsilon_N \otimes \iota_S \otimes \mu_N) \left(\overrightarrow{\text{Subject}} \otimes \overrightarrow{\text{Verb}} \otimes \overrightarrow{\text{Object}} \right) \end{aligned}$$

The simplified normal forms will become useful in practice when calculating vectors for such cases.

6 Vector Space Instantiations

In this section we demonstrate the effect of the Frobenius operations using two example instantiations. The first — which is designed perhaps

as a theoretical example rather than a suggestion for implementation — is a truth-theoretic account, similar to Coecke et al. (2010) but also allowing for degrees of truth. The second is based on the concrete implementation of Grefenstette and Sadrzadeh (2011a).

6.1 Degrees of Truth

Take N to be the vector space spanned by a set of individuals $\{\vec{n}_i\}_i$ that are mutually orthogonal. For example, \vec{n}_1 represents the individual Mary, \vec{n}_{25} represents Roger the dog, \vec{n}_{10} represents John, and so on. A sum of basis vectors in this space represents a common noun; e.g. $\overrightarrow{man} = \sum_i \vec{n}_i$, where i ranges over the basis vectors denoting men. We take S to be the one dimensional space spanned by the single vector $\vec{1}$. The unit vector spanning S represents truth value 1, the zero vector represents truth value 0, and the intermediate vectors represent degrees of truth.

A transitive verb w , which is a vector in the space $N \otimes S \otimes N$, is represented as follows:

$$\vec{w} := \sum_{ij} \vec{n}_i \otimes (\alpha_{ij} \vec{1}) \otimes \vec{n}_j$$

if \vec{n}_i w's \vec{n}_j with degree α_{ij} , for all i, j .

Further, since S is one-dimensional with its only basis vector being $\vec{1}$, the transitive verb can be represented by the following element of $N \otimes N$:

$$\sum_{ij} \alpha_{ij} \vec{n}_i \otimes \vec{n}_j \quad \text{if } \vec{n}_i \text{ w's } \vec{n}_j \text{ with degree } \alpha_{ij}$$

Restricting to either $\alpha_{ij} = 1$ or $\alpha_{ij} = 0$ provides a 0/1 meaning, i.e. either \vec{n}_i w's \vec{n}_j or not. Letting α_{ij} range over the interval $[0, 1]$ enables us to represent degrees as well as limiting cases of truth and falsity. For example, the verb “love”, denoted by \overrightarrow{love} , is represented by:

$$\sum_{ij} \alpha_{ij} \vec{n}_i \otimes \vec{n}_j \quad \text{if } \vec{n}_i \text{ loves } \vec{n}_j \text{ with degree } \alpha_{ij}$$

If we take α_{ij} to be 1 or 0, from the above we obtain the following:

$$\sum_{(i,j) \in R_{love}} \vec{n}_i \otimes \vec{n}_j$$

where R_{love} is the set of all pairs (i, j) such that \vec{n}_i loves \vec{n}_j .

Note that, with this definition, the sentence space has already been discarded, and so for this

$$\begin{aligned} \overrightarrow{\text{Subject who Verb Object}} &:= \\ (\mu_N \otimes \epsilon_N) \left(\overrightarrow{\text{Subject}} \otimes \overrightarrow{\text{Verb}} \otimes \overrightarrow{\text{Object}} \right) &= \\ (\mu_N \otimes \epsilon_N) \left(\sum_{k \in K} \vec{n}_k \otimes \left(\sum_{ij} \alpha_{ij} \vec{n}_i \otimes \vec{n}_j \right) \otimes \sum_{l \in L} \vec{n}_l \right) &= \\ = \sum_{ij, k \in K, l \in L} \alpha_{ij} \mu_N(\vec{n}_k \otimes \vec{n}_i) \otimes \epsilon_N(\vec{n}_j \otimes \vec{n}_l) &= \\ = \sum_{ij, k \in K, l \in L} \alpha_{ij} \delta_{ki} \vec{n}_i \delta_{jl} &= \\ = \sum_{k \in K, l \in L} \alpha_{kl} \vec{n}_k & \end{aligned}$$

Figure 1: Meaning computation with a subject relative pronoun

instantiation the ι map, which is the part of the relative pronoun interpretation designed to discard the relative clause after it has acted on the head noun, is not required.

For common nouns $\overrightarrow{\text{Subject}} = \sum_{k \in K} \vec{n}_k$ and $\overrightarrow{\text{Object}} = \sum_{l \in L} \vec{n}_l$, where k and l range over the sets of basis vectors representing the respective common nouns, the truth-theoretic meaning of a noun phrase modified by a subject relative clause is computed as in Figure 1. The result is highly intuitive, namely the sum of the subject individuals weighted by the degree with which they have acted on the object individuals via the verb. A similar computation, with the difference that the μ and ϵ maps are swapped, provides the truth-theoretic semantics of the object relative clause:

$$\sum_{k \in K, l \in L} \alpha_{kl} \vec{n}_l$$

The calculation and final outcome is best understood with an example.

Now only consider truth values 0 and 1. Consider the noun phrase with object relative clause “men whom Mary loves” and take N to be the vector space spanned by the set of all people; then the males form a subspace of this space, where the basis vectors of this subspace, i.e. men, are denoted by \vec{m}_l , where l ranges over the set of men which we denote by M . We set “Mary” to be the individual \vec{f}_1 , “men” to be the common noun $\sum_{l \in M} \vec{m}_l$,

$$\begin{aligned}
& \overrightarrow{\text{men whom Mary loves}} := \\
& (\epsilon_N \otimes \mu_N) \left(\vec{f}_1 \otimes \left(\sum_{(i,j) \in R_{\text{love}}} \vec{f}_i \otimes \vec{m}_j \right) \otimes \sum_{l \in M} \vec{m}_l \right) \\
& = \sum_{l \in M, (i,j) \in R_{\text{love}}} \epsilon_N(\vec{f}_1 \otimes \vec{f}_i) \otimes \mu(\vec{m}_j \otimes \vec{m}_l) \\
& = \sum_{l \in M, (i,j) \in R_{\text{love}}} \delta_{1i} \delta_{jl} \vec{m}_j \\
& = \sum_{(1,j) \in R_{\text{love}} | j \in M} \vec{m}_j
\end{aligned}$$

Figure 2: Meaning computation for example object relative clause

and “love” to be as follows:

$$\sum_{(i,j) \in R_{\text{love}}} \vec{f}_i \otimes \vec{m}_j$$

The vector corresponding to the meaning of “men whom Mary loves” is computed as in Figure 2. The result is the sum of the men basis vectors which are also loved by Mary.

The second example involves degrees of truth. Suppose we have two females Mary \vec{f}_1 and Jane \vec{f}_2 and four men $\vec{m}_1, \vec{m}_2, \vec{m}_3, \vec{m}_4$. Mary loves \vec{m}_1 with degree 1/4 and \vec{m}_2 with degree 1/2; Jane loves \vec{m}_3 with degree 1/5; and \vec{m}_4 is not loved. In this situation, we have:

$$R_{\text{love}} = \{(1, 1), (1, 2), (2, 3)\}$$

and the verb love is represented by:

$$1/4(\vec{f}_1 \otimes \vec{m}_1) + 1/2(\vec{f}_1 \otimes \vec{m}_2) + 1/5(\vec{f}_2 \otimes \vec{m}_3)$$

The meaning of “men whom Mary loves” is computed by substituting an $\alpha_{1,j}$ in the last line of Figure 2, resulting in the men whom Mary loves together with the degrees that she loves them:

$$\sum_{(1,j) \in R_{\text{love}} | j \in M} \alpha_{1,j} \vec{m}_j = 1/4 \vec{m}_1 + 1/2 \vec{m}_2$$

“men whom women love” is computed as fol-

lows, where W is the set of women:

$$\begin{aligned}
& \sum_{k \in W, l \in M, (i,j) \in R_{\text{love}}} \alpha_{ij} \epsilon_N(\vec{f}_k \otimes \vec{f}_i) \otimes \mu(\vec{m}_j \otimes \vec{m}_l) \\
& = \sum_{k \in W, l \in M, (i,j) \in R_{\text{love}}} \alpha_{ij} \delta_{ki} \delta_{jl} \vec{m}_j \\
& = \sum_{(i,j) \in R_{\text{love}} | i \in W, j \in M} \alpha_{ij} \vec{m}_j \\
& = 1/4 \vec{m}_1 + 1/2 \vec{m}_2 + 1/5 \vec{m}_3
\end{aligned}$$


The result is the men loved by Mary or Jane together with the degrees to which they are loved.

6.2 A Concrete Instantiation


In the model of Grefenstette and Sadrzadeh (2011a), the meaning of a verb is taken to be “the degree to which the verb relates properties of its subjects to properties of its object”. Clark (2013) provides some examples showing how this is an intuitive definition for a transitive verb in the categorical framework. This degree is computed by forming the sum of the tensor products of the subjects and objects of the verb across a corpus, where w ranges over instances of the verb:

$$\overline{\text{verb}} = \sum_w (\overline{\text{sbj}} \otimes \overline{\text{obj}})_w$$

Denote the vector space of nouns by N ; the above is a matrix in $N \otimes N$, depicted by a two-legged triangle as follows:



The verbs of this model do not have a sentence dimension; hence no information needs to be discarded when they are used in our setting, and so no ι map appears in the diagram of the relative clause. Inserting the above diagram in the diagrams of the normal forms results in the following for the subject relative clause (the object case is similar):



The abstract vectors corresponding to such diagrams are similar to the truth-theoretic case, with the difference that the vectors are populated from corpora and the scalar weights for noun vectors

are not necessarily 1 or 0. For subject and object noun context vectors computed from a corpus as follows:

$$\overrightarrow{\text{Subject}} = \sum_k C_k \vec{n}_k \quad \overrightarrow{\text{Object}} = \sum_l C_l \vec{n}_l$$

and the verb a linear map:

$$\overrightarrow{\text{Verb}} = \sum_{ij} C_{ij} \vec{n}_i \otimes \vec{n}_j$$

computed as above, the concrete meaning of a noun phrase modified by a subject relative clause is as follows:

$$\begin{aligned} & \sum_{kijl} C_k C_{ij} C_l \mu_N(\vec{n}_k \otimes \vec{n}_i) \epsilon_N(\vec{n}_j \otimes \vec{n}_l) \\ &= \sum_{kijl} C_k C_{ij} C_l \delta_{ki} \vec{n}_k \delta_{jl} \\ &= \sum_{kl} C_k C_{kl} C_l \vec{n}_k \end{aligned}$$

Comparing this to the truth-theoretic case, we see that the previous α_{kl} are now obtained from a corpus and instantiated to $C_k C_{kl} C_l$. To see how the above expression represents the meaning of the noun phrase, decompose it into the following:

$$\sum_k C_k \vec{n}_k \odot \sum_{kl} C_{kl} C_l \vec{n}_l$$

Note that the second term of the above, which is the application of the verb to the object, modifies the subject via point-wise multiplication. A similar result arises for the object relative clause case.

As an example, suppose that N has two dimensions with basis vectors \vec{n}_1 and \vec{n}_2 , and consider the noun phrase “dog that bites men”. Define the vectors of “dog” and “men” as follows:

$$\overrightarrow{\text{dog}} = d_1 \vec{n}_1 + d_2 \vec{n}_2 \quad \overrightarrow{\text{men}} = m_1 \vec{n}_1 + m_2 \vec{n}_2$$

and the matrix of “bites” by:

$$b_{11} \vec{n}_1 \otimes \vec{n}_2 + b_{12} \vec{n}_1 \otimes \vec{n}_2 + b_{21} \vec{n}_2 \otimes \vec{n}_1 + b_{22} \vec{n}_2 \otimes \vec{n}_2$$

Then the meaning of the noun phrase becomes:

$$\overrightarrow{\text{dog that bites men}} :=$$

$$\begin{aligned} & d_1 b_{11} m_1 \vec{n}_1 + d_1 b_{12} m_2 \vec{n}_1 + d_2 b_{21} m_1 \vec{n}_2 \\ & + d_2 b_{22} m_2 \vec{n}_2 = (d_1 \vec{n}_1 + d_2 \vec{n}_2) \odot \\ & ((b_{11} m_1 + b_{12} m_2) \vec{n}_1 + (b_{21} m_1 + b_{22} m_2) \vec{n}_2) \end{aligned}$$

Using matrix notation, we can decompose the second term further, from which the application of the verb to the object becomes apparent:

$$\begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \times \begin{pmatrix} m_1 \\ m_2 \end{pmatrix}$$

Hence for the whole clause we obtain:

$$\overrightarrow{\text{dog}} \odot (\overrightarrow{\text{bites}} \times \overrightarrow{\text{men}})$$


Again this result is highly intuitive: assuming that the basis vectors of the noun space represent properties of nouns, the meaning of “dog that bites men” is a vector representing the properties of dogs, which have been modified (via multiplication) by those properties of individuals which bite men. Put another way, those properties of dogs which overlap with properties of biting things get accentuated.

7 Conclusion and Future Directions


In this paper, we have extended the compact categorical semantics of Coecke et al. (2010) to analyse meanings of relative clauses in English from a vector space point of view. The resulting vector space semantics of the pronouns and clauses is based on the Frobenius algebraic operations on vector spaces: they reveal the internal structure, or what we call *anatomy*, of the relative clauses.

The methodology pursued in this paper and the Frobenius operations can be used to provide semantics for other relative pronouns and also other closed-class words such as prepositions and determiners. In each case, the grammatical type of the word and a detailed analysis of the role of these words in the meaning of the phrases in which they occur would be needed. In some cases, it may be necessary to introduce a linear map to represent the meaning of the word, for instance to distinguish the preposition *on* from *in*.

The contribution of this paper is best demonstrated via the string diagrammatic representations of the vector space meanings of these clauses. A noun phrase modified by a subject relative clause, which before this paper was depicted as follows:



will now include the internal anatomy of its relative pronoun:



This internal structure shows how the information from the noun flows through the relative pronoun to the rest of the clause and how it interacts with the other words. We have instantiated this vector space semantics using truth-theoretic and corpus-based examples.

One aspect of our example spaces which means that they work particularly well is that the sentence dimension in the verb is already discarded, which means that the ι maps are not required (as discussed above). Another feature is that the simple nature of the models means that the μ map does not lose any information, even though it takes the diagonal of a matrix and hence in general throws information away. The effect of the ι and μ maps in more complex representations of the verb remains to be studied in future work.

On the practical side, what we offer in this paper is a method for building appropriate vector representations for relative clauses. As a result, when presented with a relative clause, we are able to build a vector for it, only by relying on the vector representations of the words in the clause and the grammatical role of the relative pronoun. We do not need to retrieve information from a corpus to be able to build a vector or linear map for the relative pronoun, neither will we end up having to discard the pronoun and ignore the role that it plays in the meaning of the clause (which was perhaps the best option available before this paper). However, the Frobenius approach and our claim that the resulting vectors are ‘appropriate’ requires an empirical evaluation. Tasks such as the term definition task from Kartsaklis et al. (2013) (which also uses Frobenius algebras but for a different purpose) are an obvious place to start. More generally, the subfield of compositional distributional semantics is a growing and active one (Mitchell and Lapata, 2008; Baroni and Zamparelli, 2010; Zanzotto et al., 2010; Socher et al., 2011), for which we argue that high-level mathematical investigations such as this paper, and also Clarke (2008), can play a crucial role.

Acknowledgements

We would like to thank Dimitri Kartsaklis and Laura Rimell for helpful comments. Stephen Clark was supported by ERC Starting Grant Dis-CoTex (30692). Bob Coecke and Stephen Clark are supported by EPSRC Grant EP/I037512/1. Mehrnoosh Sadrzadeh is supported by an EPSRC CAF EP/J002607/1.

References

- J.C. Baez and J. Dolan. 1995. Higher-dimensional algebra and topological quantum field theory. *Journal of Mathematical Physics*, 36:6073–6105.
- M. Baroni and R. Zamparelli. 2010. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Conference on Empirical Methods in Natural Language Processing (EMNLP-10)*, Cambridge, MA.
- A. Carboni and R. F. C. Walters. 1987. Cartesian bicategories. I. *J. Pure and Applied Algebra*, 49:11–32.
- S. Clark. 2013. Type-driven syntax and semantics for composing meaning vectors. In Chris Heunen, Mehrnoosh Sadrzadeh, and Edward Grefenstette, editors, *Quantum Physics and Linguistics: A Compositional, Diagrammatic Discourse*, pages 359–377. Oxford University Press.
- D. Clarke. 2008. *Context-theoretic Semantics for Natural Language: An Algebraic Framework*. Ph.D. thesis, University of Sussex.
- B. Coecke and E. Paquette. 2008. Introducing categories to the practicing physicist. In B. Coecke, editor, *New Structures for Physics*, volume 813 of *Lecture Notes in Physics*, pages 167–271. Springer.
- B. Coecke, D. Pavlovic, and J. Vicary. 2008. A new description of orthogonal bases. *Mathematical Structures in Computer Science*, 1:269–272.
- B. Coecke, M. Sadrzadeh, and S. Clark. 2010. Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis*, 36:345–384.
- F. G. Frobenius. 1903. *Theorie der hyperkomplexen Größen*. Preussische Akademie der Wissenschaften Berlin: Sitzungsberichte der Preußischen Akademie der Wissenschaften zu Berlin. Reichsdr.
- E. Grefenstette and M. Sadrzadeh. 2011a. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1394–1404.
- E. Grefenstette and M. Sadrzadeh. 2011b. Experimenting with transitive verbs in a discocat. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics (GEMS)*.

- A. Joyal and R. Street. 1991. The Geometry of Tensor Calculus, I. *Advances in Mathematics*, 88:55–112.
- D. Kartsaklis, M. Sadrzadeh, S. Pulman, and B. Coecke. 2013. Reasoning about meaning in natural language with compact closed categories and Frobenius algebras. In J. Chubb, A. Eskandarian, and V. Harizanov, editors, *Logic and Algebraic Structures in Quantum Computing and Information*, Association for Symbolic Logic Lecture Notes in Logic. Cambridge University Press.
- G. M. Kelly and M. L. Laplaza. 1980. Coherence for compact closed categories. *Journal of Pure and Applied Algebra*, 19:193–213.
- J. Kock. 2003. *Frobenius algebras and 2D topological quantum field theories*, volume 59 of *London Mathematical Society student texts*. Cambridge University Press.
- J. Lambek. 1958. The Mathematics of Sentence Structure. *American Mathematics Monthly*, 65:154–170.
- J. Lambek. 1999. Type Grammar Revisited Logical Aspects of Computational Linguistics. In *Logical Aspects of Computational Linguistics*, volume 1582 of *Lecture Notes in Computer Science*, pages 1–27. Springer Berlin / Heidelberg.
- J. Mitchell and M. Lapata. 2008. Vector-based models of semantic composition. In *Proceedings of ACL-08*, pages 236–244, Columbus, OH.
- R. Montague. 1974. English as a formal language. In R. H. Thomason, editor, *Formal philosophy: Selected Papers of Richard Montague*, pages 189–223. Yale University Press.
- A. Preller and J. Lambek. 2007. Free compact 2-categories. *Mathematical Structures in Computer Science*, 17:309–340.
- H. Schütze. 1998. Automatic word sense discrimination. *Computational Linguistics*, 24:97–123.
- R. Socher, J. Pennington, E. Huang, A. Y. Ng, and C. D. Manning. 2011. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK.
- F. M. Zanzotto, I. Korkontzelos, F. Fallucchi, and S. Manandhar. 2010. Estimating linear models for compositional distributional semantics. In *Proceedings of COLING*, Beijing, China.