



HAL
open science

Distributional Sentence Entailment Using Density Matrices

Esma Balkir, Mehrnoosh Sadrzadeh, Bob Coecke

► **To cite this version:**

Esma Balkir, Mehrnoosh Sadrzadeh, Bob Coecke. Distributional Sentence Entailment Using Density Matrices. 1st International Conference on Theoretical Computer Science (TTCS), Aug 2015, Tehran, Iran. pp.1-22, 10.1007/978-3-319-28678-5_1 . hal-01446260

HAL Id: hal-01446260

<https://inria.hal.science/hal-01446260>

Submitted on 25 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Distributional Sentence Entailment Using Density Matrices

Esma Balkir¹, Mehrnoosh Sadrzadeh¹, and Bob Coecke²

¹ Queen Mary University of London

² University of Oxford

Abstract. Categorical compositional distributional model of Clark, Coecke, and Sadrzadeh suggests a way to combine grammatical composition of the formal, type logical models with the corpus based, empirical word representations of distributional semantics. This paper contributes to the project by expanding the model to also capture entailment relations. This is achieved by extending the representations of words from points in meaning space to density operators, which are probability distributions on the subspaces of the space. A symmetric measure of similarity and an asymmetric measure of entailment is defined, where lexical entailment is measured using von Neumann entropy, the quantum variant of Kullback-Leibler divergence. Lexical entailment, combined with the composition map on word representations, provides a method to obtain entailment relations on the level of sentences. Truth theoretic and corpus-based examples are provided.

1 Introduction

The term *distributional semantics* is almost synonymous with the term *vector space models of meaning*. This is because vector spaces are natural candidates for modelling the *distributional hypothesis* and contextual similarity between words [11]. In a nutshell, this hypothesis says that words that often occur in the same contexts have similar meanings. So for instance, ‘ale’ and ‘lager’ are similar since they both often occur in the context of ‘beer’, ‘pub’, and ‘pint’. The obvious practicality of these models, however, does not guarantee that they possess the expressive power needed to model all aspects of meaning. Current distributional models mostly fall short of successfully modelling subsumption and entailment [19]. There are a number of models that use distributional similarity to enhance textual entailment [13, 4]. However, most of the work from the distributional semantics community has been focused on developing more sophisticated metrics on vector representations [20, 17, 27].

In this paper we suggest the use of density matrices instead of vector spaces as the basic distributional representations for the meanings of words. Density matrices are widely used in quantum mechanics, and are a generalization of vectors. There are several advantages to using density matrices to model meaning. Firstly, density matrices have the expressive power to represent all the information vectors can represent: they are a suitable implementation of the distributional hypothesis. They come equipped with a measure of information content, and so provide a natural way of implementing asymmetric relations between words such as hyponymy-hypernymy

relations. Furthermore, they form a compact closed category. This allows the previous work of [7, 10] on obtaining representations for meanings of sentences from the meaning of words to be applicable to density matrices. The categorical map from meanings of words to the meaning of the sentence respects the order induced by the relative entropy of density matrices. This promises, given suitable representations of individual words, a method to obtain entailment relations on the level of sentences, inline with the lexical entailment of *natural logic*, e.g., see [21], rather than the traditional logical entailment of Montague semantics.

Related Work. This work builds upon and relates to the literature on compositional distributional models, distributional lexical entailment, and the use of density matrices in computational linguistics and information retrieval.

There has been a recent interest in methods of composition within the distributional semantics framework. There are a number of composition methods in literature. See [14] for a survey of compositional distributional models and a discussion of their strengths and weaknesses. This work extends the work presented in [7, 10], a compositional model based on category theory. Their model was shown to outperform the competing compositional models in [15].

Research on distributional entailment has mostly been focused on lexical entailment. One notable exception is the work in [3], which uses the distributional data on adjective-noun and quantifier-noun pairs to train a classifier; the results are then utilized to detect novel noun pairs that have the same relation. There are a number of non-symmetric lexical entailment measures, e.g., see [27, 8, 17, 19], all of which rely on some variation of the *Distributional Inclusion Hypothesis*: “If u is semantically narrower than v , then a significant number of salient distributional features of u are also included in the feature vector of v ” [17]. In their experiments, the authors of [12] show that while if a word v entails another word w then the characteristic features of v is a subset of the ones for w , it is not necessarily the case that the inclusion of the characteristic features v in w indicate that v entails w . One of their suggestions for increasing the prediction power of their method is to include more than one word in the features.

The work presented in [23] uses a measure based on entropy to detect hyponym-hypernym relationships in given pairs. The measure they suggest rely on the hypothesis that hypernyms are semantically more general than hyponyms, and therefore tend to occur in less informative contexts. The authors of [16] rely on a very similar idea, and use KL-divergence between the target word and the basis words to quantify the semantic content of the target word. They conclude that this method performs equally well in detecting hyponym-hypernym pairs as their baseline prediction method that only considers the overall frequency of the word in corpus. They reject the hypothesis that more general words occur in less informative contexts. Their method differs from ours in that they use relative entropy to quantify the overall information content of a word, and not to compare two target words to each other.

The work presented in [22] extend the compositional model of [7, 10] to include density matrices as we do, but use it for modeling homonymy and polysemy. Their approach is complementary to ours, and in fact, they show that it is possible to merge the two constructions. The work presented in [6] uses density matrices to model context effects in a conceptual space. In their quantum mechanics inspired model, words are

represented by mixed states and each eigenstate represents a sense of the word. Context effects are then modelled as quantum collapse. The authors in [5] use density matrices to encode dependency neighbourhoods, with the aim of modelling context effects in similarity tasks; the work presented in [26] uses density matrices to sketch out a theory of information retrieval, and connects the logic of the space and of density matrices via an order relation that makes the set of projectors in a Hilbert space into a complete lattice. They then use this order to define an entailment relation. Finally, the work presented in [25] shows that using density matrices to represent documents provides significant improvement on realistic IR tasks.

This paper is based on the MSc Thesis of the first author [2].


2 Background

Definition 1. A monoidal category is **compact closed** if for any object A , there are left and right dual objects, i.e. objects A^r and A^l , and morphisms $\eta^l : I \rightarrow A \otimes A^l$, $\eta^r : I \rightarrow A^r \otimes A$, $\epsilon^l : A^l \otimes A \rightarrow I$ and $\epsilon^r : A \otimes A^r \rightarrow I$ that satisfy:


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

Compact closed categories are used to represent *correlations*, and in categorical quantum mechanics they model maximally entangled states. [1] The η and ϵ maps are useful in modeling the interactions of the different parts of a system. To see how this relates to natural language, consider a simple sentence with an object, a subject and a transitive verb. The meaning of the entire sentence is not simply an accumulation of the meanings of its individual words, but depends on how the transitive verb relates the subject and the object. The η and ϵ maps provide the mathematical formalism to specify such interactions. The distinct left and right duals ensure that compact closed categories can take word order into account.


There is a graphical calculus used to reason about monoidal categories [9]. In the graphical language, objects are wires, and morphisms are boxes with incoming and outgoing wires of types corresponding to the input and output types of the morphism. The identity object is depicted as empty space, so a state $\psi : I \rightarrow A$ is depicted as a box with no input wire and an output wire with type A . The duals of states are called *effects*, and they are of type $A \rightarrow I$. Let $f : A \rightarrow B$, $g : B \rightarrow C$ and $h : C \rightarrow D$, and $1_A : A \rightarrow A$ the identity function on A . 1_A , f , $f \otimes h$, $g \circ f$ are depicted as follows:




The state $\psi : I \rightarrow A$, the effect $\pi : A \rightarrow I$, and the scalar $\psi \circ \pi$ are depicted as follows:




The maps $\eta^l, \eta^r, \epsilon^l$ and ϵ^r take the following forms in the graphical calculus:



The axioms of compact closure, referred to as the *snake identities* because of the visual form they take in the graphical calculus, are represented as follows:



More generally, the reduction rules for diagrammatic calculus allow continuous deformations. One such deformation we will make use of is the *swing rule*:



Definition 2. [18] A **pregroup** $(P, \leq, \cdot, 1, (-)^l, (-)^r)$ is a partially ordered monoid in which each element a has both a left adjoint a^l and a right adjoint a^r such that $a^l a \leq 1 \leq a a^l$ and $a a^r \leq 1 \leq a^r a$.

If $a \leq b$ it is common practice to write $a \rightarrow b$ and say that a reduces to b . This terminology is useful when pregroups are applied to natural language, where each word gets assigned a pregroup type freely generated from a set of basic elements. The sentence is deemed to be grammatical if the concatenation of the types of the words reduce to the simple type of a sentence. For example reduction for a simple transitive sentence is $n(n^r s n^l)n \rightarrow 1 s n^l n \rightarrow 1 s 1 \rightarrow s$.

A pregroup \mathbf{P} is a concrete instance of a compact closed category. The $\eta^l, \eta^r, \epsilon^l, \epsilon^r$ maps are $\eta^l = [1 \leq p \cdot p^l], \epsilon^l = [p^l \cdot p \leq 1], \eta^r = [1 \leq p^r \cdot p], \epsilon^r = [p \cdot p^r \leq 1]$.

FVect as a concrete compact closed category. Finite dimensional vector spaces over the base field \mathbb{R} , together with linear maps form a monoidal category, referred to as **FVect**. The monoidal tensor is the usual vector space tensor and the monoidal unit is the base field \mathbb{R} . It is also a compact closed category where $V^l = V^r = V$. The compact closed maps are defined as follows:

Given a vector space V with basis $\{\vec{e}_i\}_i$,

$$\begin{aligned} \eta_V^l = \eta_V^r : \mathbb{R} &\rightarrow V \otimes V & \epsilon_V^l = \epsilon_V^r : V \otimes V &\rightarrow \mathbb{R} \\ 1 &\mapsto \sum_i \vec{e}_i \otimes \vec{e}_i & \sum_{ij} c_{ij} \vec{v}_i \otimes \vec{w}_i &\mapsto \sum_{ij} c_{ij} \langle \vec{v}_i | \vec{w}_i \rangle \end{aligned}$$

Categorical representation of meaning space. The tensor in **FVect** is commutative up to isomorphism. This causes the left and the right adjoints to be the same, and thus the left and the right compact closed maps to coincide. Thus **FVect** by itself cannot take the effect of word ordering on meaning into account. [7, 10] propose a way around this obstacle by considering the product category **FVect** \times **P** where **P** is a pregroup.

Objects in **FVect** \times **P** are of the form (V, p) , where V is the vector space for the representation of meaning and p is the pregroup type. There exists a morphism $(f, \leq) : (V, p) \rightarrow (W, q)$ if there exists a morphism $f : V \rightarrow W$ in **FVect** and $p \leq q$ in **P**.


The compact closed structure of **FVect** and **P** lifts componentwise to the product category **FVect** \times **P**:

$$\begin{aligned} \eta^l : (\mathbb{R}, 1) &\rightarrow (V \otimes V, p \cdot p^l) & \eta^r : (\mathbb{R}, 1) &\rightarrow (V \otimes V, p^r \cdot p) \\ \epsilon^l : (V \otimes V, p^l \cdot p) &\rightarrow (\mathbb{R}, 1) & \epsilon^r : (V \otimes V, p \cdot p^r) &\rightarrow (\mathbb{R}, 1) \end{aligned}$$

Definition 3. An object (V, p) in the product category is called a **meaning space**, where V is the vector space in which the meanings $\vec{v} \in V$ of strings of type p live.

Definition 4. From-meanings-of-words-to-the-meaning-of-the-sentence map. Let $v_1 v_2 \dots v_n$ be a string of words, each v_i with a meaning space representation $\vec{v}_i \in (V_i, p_i)$. Let $x \in P$ be a pregroup type such that $[p_1 p_2 \dots p_n \leq x]$ Then the meaning vector for the string is $\vec{v}_1 \vec{v}_2 \dots \vec{v}_n := f(\vec{v}_1 \otimes \vec{v}_2 \otimes \dots \otimes \vec{v}_n) \in (W, x)$, where f is defined to be the application of the compact closed maps obtained from the reduction $[p_1 p_2 \dots p_n \leq x]$ to the composite vector space $V_1 \otimes V_2 \otimes \dots \otimes V_n$.

This framework uses the maps of the pregroup reductions and the elements of objects in **FVect**. The diagrammatic calculus provides a tool to reason about both. As an example, take the sentence ‘‘John likes Mary’’. It has the pregroup type $nn^r sn^l n$, and the vector representations $\overrightarrow{John}, \overrightarrow{Mary} \in V$ and $\overrightarrow{likes} \in V \otimes S \otimes V$. The morphism in **FVect** \times **P** corresponding to the map defined in Definition 4 is of type $(V \otimes (V \otimes S \otimes V) \otimes V, nn^r sn^l n) \rightarrow (S, s)$. From the pregroup reduction $[nn^r sn^l n \rightarrow s]$ we obtain the compact closed maps $\epsilon^r 1 \epsilon^l$. In **FVect** this translates into $\epsilon_V \otimes 1_S \otimes \epsilon_V : V \otimes (V \otimes S \otimes V) \otimes V \rightarrow S$. This map, when applied to $\overrightarrow{John} \otimes \overrightarrow{likes} \otimes \overrightarrow{Mary}$, has the following depiction in the diagrammatic calculus:




Note that this construction treats the verb ‘likes’ essentially as a relation that takes two inputs of type V , and outputs a vector of type S . For the explicit calculation, note that $\overrightarrow{\text{likes}} = \sum_{ijk} c_{ijk} \overrightarrow{v}_i \otimes \overrightarrow{s}_j \otimes \overrightarrow{v}_k$, where $\{\overrightarrow{v}_i\}_i$ is an orthonormal basis for V and $\{\overrightarrow{s}_j\}_j$ is an orthonormal basis for S . Then

$$\overrightarrow{\text{John likes Mary}} = \epsilon_V \otimes 1_S \otimes \epsilon_V (\overrightarrow{\text{John}} \otimes \overrightarrow{\text{likes}} \otimes \overrightarrow{\text{Mary}}) \quad (1)$$

$$= \sum_{ijk} \langle \overrightarrow{\text{John}} | \overrightarrow{v}_i \rangle \overrightarrow{s}_j \langle \overrightarrow{v}_k | \overrightarrow{\text{Mary}} \rangle \quad (2)$$

The reductions in diagrammatic calculus help reduce the final calculation to a simpler term. The non-reduced reduction, when expressed in dirac notation is $(\langle \epsilon_V^r | \otimes 1_S \otimes \langle \epsilon_V^l |) \circ \overrightarrow{\text{John}} \otimes \overrightarrow{\text{likes}} \otimes \overrightarrow{\text{Mary}}$. But we can swing $\overrightarrow{\text{John}}$ and $\overrightarrow{\text{Mary}}$ in accord with the reduction rules in the diagrammatic calculus. The diagram then reduces to:



This results in a simpler expression that needs to be calculated: $(\langle \overrightarrow{\text{John}} | \otimes 1_S \otimes \langle \overrightarrow{\text{Mary}} |) \circ \overrightarrow{\text{likes}}$.


3 Density Matrices as Elements of a Compact Closed Category

Recall that in \mathbf{FVect} , vectors $|\overrightarrow{v}\rangle \in V$ are in one-to-one correspondence with morphisms of type $v : I \rightarrow V$. Likewise, pure states of the form $|\overrightarrow{v}\rangle\langle\overrightarrow{v}|$ are in one-to-one correspondence with morphisms $v \circ v^\dagger : V \rightarrow V$ such that $v^\dagger \circ v = \text{id}_I$, where v^\dagger denotes the adjoint of v (notice that this corresponds to the condition that $\langle v | v \rangle = 1$). A general (mixed) state ρ is a positive morphism of the form $\rho : V \rightarrow V$. One can re-express the mixed states $\rho : V \rightarrow V$ as elements $\rho : I \rightarrow V^* \otimes V$. Here $V^* = V^l = V^r = V$.

Definition 5. f is a **completely positive map** if f is positive for any positive operator A , and $(\text{id}_V \otimes f)B$ is positive for any positive operator B and any space V .

Completely positive maps in \mathbf{FVect} form a monoidal category [24]. Thus one can define a new category $\mathbf{CPM}(\mathbf{FVect})$ where the objects of $\mathbf{CPM}(\mathbf{FVect})$ are the same as those of \mathbf{FVect} , and morphisms $A \rightarrow B$ in $\mathbf{CPM}(\mathbf{FVect})$ are completely positive maps $A^* \otimes A \rightarrow B^* \otimes B$ in \mathbf{FVect} . The elements $I \rightarrow A$ in $\mathbf{CPM}(\mathbf{FVect})$ are of the form $I^* \otimes I \rightarrow A^* \otimes A$ in \mathbf{FVect} , providing a monoidal category with density matrices as its elements.


CPM(FVect) in graphical calculus. A morphism $\rho : A \rightarrow A$ is positive if and only if there exists a map $\sqrt{\rho}$ such that $\rho = \sqrt{\rho}^\dagger \circ \sqrt{\rho}$. In **FVect**, the isomorphism between $\rho : A \rightarrow A$ and $\lceil \rho \rceil : I \rightarrow A^* \otimes A$ is provided by $\eta^l = \eta^r$. The graphical representation of ρ in **FVect** then becomes:



Notice that the categorical definition of a positive morphism coincides with the definition of a positive operator in a vector space, where $\sqrt{\rho}$ is the square root of the operator.

The graphical depiction of completely positive morphisms come from the following theorem:

Theorem 1. (*Stinespring Dilation Theorem*) $f : A^* \otimes A \rightarrow B^* \otimes B$ is completely positive if and only if there is an object C and a morphism $\sqrt{f} : A \rightarrow C \otimes B$ such that the following equation holds:




\sqrt{f} and C here are not unique. For the proof of the theorem see [24].

Theorem 2. **CPM(FVect)** is a compact closed category where as in **FVect**, $V^r = V^l = V$ and the compact closed maps are defined to be:

$$\begin{aligned} \eta^l &= (\eta_V^r \otimes \eta_V^l) \circ (\mathbf{1}_V \otimes \sigma \otimes \mathbf{1}_V) & \eta^r &= (\eta_V^l \otimes \eta_V^r) \circ (\mathbf{1}_V \otimes \sigma \otimes \mathbf{1}_V) \\ \epsilon^l &= (\mathbf{1}_V \otimes \sigma \otimes \mathbf{1}_V) \circ (\epsilon_V^r \otimes \epsilon_V^l) & \epsilon^r &= (\mathbf{1}_V \otimes \sigma \otimes \mathbf{1}_V) \circ (\epsilon_V^l \otimes \epsilon_V^r) \end{aligned}$$

where σ is the swap map defined as $\sigma(v \otimes w) = (w \otimes v)$.

Proof. The graphical construction of the compact closed maps boils down to doubling the objects and the wires. The identities are proved by adding bends in the wires. Consider the diagram for η^r :



These maps satisfy the axioms of compact closure since the components do.

The concrete compact closed maps are as follows:

$$\begin{aligned} \eta^l = \eta^r : \mathbb{R} &\rightarrow (V \otimes V) \otimes (V \otimes V) \\ &:: 1 \mapsto \sum_i \vec{e}_i \otimes \vec{e}_i \otimes \sum_j \vec{e}_j \otimes \vec{e}_j \\ \epsilon^l = \epsilon^r : (V \otimes V) \otimes (V \otimes V) &\rightarrow \mathbb{R} \\ &:: \sum_{ijkl} c_{ijkl} \vec{v}_i \otimes \vec{w}_j \otimes \vec{u}_k \otimes \vec{p}_l \mapsto \sum_{ijkl} c_{ijkl} \langle \vec{v}_i | \vec{u}_k \rangle \langle \vec{w}_j | \vec{p}_l \rangle \end{aligned}$$

Let $\rho : V_1 \otimes V_2 \otimes \dots \otimes V_n \rightarrow V_1 \otimes V_2 \otimes \dots \otimes V_n$ be a density operator defined on an arbitrary composite space $V_1 \otimes V_2 \otimes \dots \otimes V_n$. Then it has the density matrix representation $\rho : I \rightarrow (V_1 \otimes V_2 \otimes \dots \otimes V_n)^* \otimes (V_1 \otimes V_2 \otimes \dots \otimes V_n)$. Since the underlying category \mathbf{FVect} is symmetric, it has the swap map σ . This provides us with the isomorphism:

$$(V_1 \otimes V_2 \otimes \dots \otimes V_n)^* \otimes (V_1 \otimes V_2 \otimes \dots \otimes V_n) \sim (V_1^* \otimes V_1) \otimes (V_2^* \otimes V_2) \otimes \dots \otimes (V_n^* \otimes V_n)$$

So ρ can be equivalently expressed as $\rho : I \rightarrow (V_1^* \otimes V_1) \otimes (V_2^* \otimes V_2) \otimes \dots \otimes (V_n^* \otimes V_n)$. With this addition, we can simplify the diagrams used to express density matrices by using a single thick wire for the doubled wires. Doubled compact closed maps can likewise be expressed by a single thick wire.

$$\begin{array}{ccc} \downarrow \text{ (thick)} & := & \downarrow \downarrow \text{ (thin)} \\ \curvearrowright \text{ (thick)} & := & \curvearrowright \curvearrowright \text{ (thin)} \\ \curvearrowleft \text{ (thick)} & := & \curvearrowleft \curvearrowleft \text{ (thin)} \end{array}$$

The diagrammatic expression of a from-meanings-of-words-to-the-meaning-of-the-sentence map using density matrices will therefore look exactly like the depiction of it in \mathbf{FVect} , but with thick wires.

4 Using Density Matrices to Model Meaning

If one wants to use the full power of density matrices in modelling meaning, one needs to establish an interpretation for the distinction between *mixing* and *superposition* in the context of linguistics. Let *contextual features* be the salient, quantifiable features of the contexts a word is observed in. Let the basis of the space be indexed by such contextual features. Individual *contexts*, such as words in an n -word window of a text, can be represented as the superposition of the bases corresponding to the contextual features observed in it. So each context corresponds to a pure state. Words are then probability distributions over the contexts they appear in. The simple co-occurrence model can be cast as a special case of this more general approach, where features and contexts are the same. Then all word meanings are mixtures of basis vectors, and they all commute with each other.

Similarity for density matrices. Fidelity is a good measure of similarity between two density matrix representations of meaning because of its properties listed below.

Definition 6. *The fidelity of two density operators ρ and σ is $F(\rho, \sigma) := \text{tr}\sqrt{\rho^{1/2}\sigma\rho^{1/2}}$.*

Some useful properties of fidelity are:

1. $F(\rho, \sigma) = F(\sigma, \rho)$
2. $0 \leq F(\rho, \sigma) \leq 1$
3. $F(\rho, \sigma) = 1$ if and only if $\rho = \sigma$
4. If $|\phi\rangle\langle\phi|$ and $|\psi\rangle\langle\psi|$ are two pure states, their fidelity is equal to $|\langle\phi|\psi\rangle|$.

These properties ensure that if the representations of two words are not equal to each other, they will not be judged perfectly similar, and, if two words are represented as projections onto one dimensional subspaces, their similarity value will be equal to the usual cosine similarity of the vectors.

Entailment for density matrices. To develop a theory of entailment using density matrices as the basic representations, we assume the following hypothesis:

Definition 7 (Distributional Hypothesis for Hyponymy). *The meaning of a word w subsumes the meaning of a word v if and only if it is appropriate to use w in all the contexts v is used.*

This is a slightly more general version of the *Distributional Inclusion Hypothesis* (DIH) stated in [17]. The difference lies in the additional power the density matrix formalism provides: the distinction between mixing and superposition. Further, DIH only considers whether or not the target word occurs together with the salient distributional feature at all, and ignores any possible statistically significant correlations of features; here again, the density matrix formalism offers a solution.

Note that [12] show that while there is ample evidence for the distributional inclusion hypothesis, this in itself does not necessarily provide a method to detect hyponymy-hypernymy pairs. One of their suggestions for improvement is to consider more than one word in the features, equivalent to what we do here by taking correlations into account in a co-occurrence space where the bases are context words.

Relative entropy quantifies the distinguishability of one distribution from another. The idea of using relative entropy to model hyponymy is based on the assumption that the distinguishability of one word from another given its usual contexts provides us with a good metric for hyponymy. For example, if one is given a sentence with the word *dog* crossed out, it will be not be possible for sure to know whether the crossed out word is not *animal* just from the context (except perhaps very particular decelerational sentences which rely on world knowledge, such as ‘All – bark’.)

Definition 8. *The (quantum) relative entropy of two density matrices ρ and σ is $N(\rho||\sigma) := \text{tr}(\rho \log \rho) - \text{tr}(\rho \log \sigma)$, where $0 \log 0 = 0$ and $x \log 0 = \infty$ when $x \neq 0$ by convention.*

Definition 9. *The representativeness between ρ and σ is $R(\rho, \sigma) := 1/(1 + N(\rho||\sigma))$, where $N(\rho||\sigma)$ is the quantum relative entropy between ρ and σ .*

Quantum relative entropy is always non-negative. For two density matrices ρ and σ , $N(\rho||\sigma) = \infty$ if $\text{supp}(\rho) \cap \ker(\sigma) \neq 0$, and is finite otherwise. The following is a direct consequence of these properties:

Corollary 1. *For all density matrices ρ and σ , $R(\rho, \sigma) \leq 1$ with equality if and only if $\rho = \sigma$, and $0 \leq R(\rho, \sigma)$ with equality if and only if $\text{supp}(\rho) \cap \ker(\sigma) \neq 0$*

The second part of the corollary reflects the idea that if there is a context in which it is appropriate to use v but not w , then v is perfectly distinguishable from w . Such contexts are exactly those that fall within $\text{supp}(\rho) \cap \ker(\sigma)$.

Characterizing hyponyms. The quantitative measure on density matrices given by representativeness provide a qualitative preorder on meaning representations as follows:

$$\begin{aligned} \rho \prec \sigma & \text{ if } R(\rho, \sigma) > 0 \\ \rho \sim \sigma & \text{ if } \prec \sigma \text{ and } \sigma \prec \rho \end{aligned}$$

Proposition 1. *The following are equivalent:*

1. $\rho \prec \sigma$
2. $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$
3. *There exists a positive operator ρ' and $p > 0$ such that $\sigma = p\rho + \rho'$*

Proof. (1) \Rightarrow (2) and (2) \Rightarrow (1) follow directly from Corollary 1.

(2) \Rightarrow (3) since $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$ implies that there exists a $p > 0$ such that $\sigma - p\rho$ is positive. Setting $\rho' = \sigma - p\rho$ gives the desired equality.

(3) \Rightarrow (2) since $p > 0$, and so $\text{supp}(\rho) \subseteq \text{supp}(\sigma) = \text{supp}(p\rho + \rho')$.

The equivalence relation \sim groups any two density matrices ρ and σ with $\text{supp}(\rho) = \text{supp}(\sigma)$ into the same equivalence class, thus maps the set of density matrices on a Hilbert space \mathcal{H} onto the set of projections on \mathcal{H} . The projections are in one-to-one correspondence with the subspaces of \mathcal{H} and they form an orthomodular lattice, providing a link to the logical structure of the Hilbert space [26] aims to exploit by using density matrices in IR.

Let \hat{w} and \hat{v} be density matrix representations of the words v and w . Then v is a hyponym of w in this model if $\hat{v} \prec \hat{w}$ and $\hat{v} \sim \hat{w}$.

Notice that even though this ordering on density matrices extracts a *yes/no* answer for the question “is v a hyponym of w ?”, the existence of the quantitative measure lets us to also quantify the extent to which v is a hyponym of w . This provides some flexibility in characterizing hyponymy through density matrices in practice. Instead of calling v a hyponym of w even when $R(\hat{v}, \hat{w})$ gets arbitrarily small, one can require the representativeness to be above a certain threshold ϵ . This modification, however, has the down side of causing the transitivity of hyponymy to fail.

5 From meanings of words to the meanings of sentences passage

As in the case for $\mathbf{FVect} \times \mathbf{P}$, $\mathbf{CPM}(\mathbf{FVect}) \times \mathbf{P}$ is a compact closed category, where the compact closed maps of $\mathbf{CPM}(\mathbf{FVect})$ and \mathbf{P} lift component-wise to the product category.

Definition 10. A *meaning space* in this new category is a pair $(V^* \otimes V, p)$ where $V^* \otimes V$ is the space in which density matrices $v : I \rightarrow V^* \otimes V$ of the pregroup type p live.


Definition 11. Let $v_1 v_2 \dots v_n$ be a string of words, each v_i with a meaning space representation $\widehat{v}_i \in (V_i^* \otimes V_i, p_i)$. Let $x \in P$ be a pregroup type such that $[p_1 p_2 \dots p_n \leq x]$. Then the meaning density matrix for the string is defined as:

$$v_1 \widehat{v_2 \dots v_n} := f(\widehat{v}_1 \otimes \widehat{v}_2 \otimes \dots \otimes \widehat{v}_n) \in (W^* \otimes W, x)$$

where f is defined to be the application of the compact closed maps obtained from the reduction $[p_1 p_2 \dots p_n \leq x]$ to the composite density matrix space $(V_1 \otimes V_1^*) \otimes (V_2^* \otimes V_2) \otimes \dots \otimes (V_n^* \otimes V_n)$.

From a high level perspective, the reduction diagrams for $\mathbf{CPM}(\mathbf{FVect}) \times \mathbf{P}$ look no different than the original diagrams for $\mathbf{FVect} \times \mathbf{P}$, except that we depict them with thick instead of thin wires. Consider the previous example: ‘‘John likes Mary’’. It has the pregroup type $n(n^r sn^l)n$, and the compact closed maps obtained from the pregroup reduction is $(\epsilon^r \otimes 1 \otimes \epsilon^l)$.

One can also depict the diagram together with the internal anatomy of the density representations in \mathbf{FVect} :



The graphical reductions for compact closed categories can be applied to the diagram, establishing $(\epsilon^r \otimes 1 \otimes \epsilon^l)(\widehat{John} \otimes \widehat{likes} \otimes \widehat{Mary}) = (\widehat{John} \otimes 1 \otimes \widehat{Mary}) \circ \widehat{likes}$.

As formalised in natural logic, one expects that if the subject and object of a sentence are common nouns which are, together with the verb of the sentence, moreover, upward monotone, then if these are replaced by their hyponyms, then the meanings of the original and the modified sentences would preserve this hyponymy. The following proposition shows that the sentence meaning map for simple transitive sentences achieves exactly that.

Theorem 3. If $\rho, \sigma, \delta, \gamma \in (N^* \otimes N, n)$, $\alpha, \beta \in (N^* \otimes N \otimes S^* \otimes S \otimes N^* \otimes N, n^l sn^r)$ then $\rho \prec \sigma$, $\delta \prec \gamma$ and $\alpha \prec \beta$ then

$$f(\rho \otimes \alpha \otimes \delta) \prec f(\sigma \otimes \beta \otimes \gamma)$$

where f is the from-meanings-of-words-to-the-meaning-of-the-sentence map in definition 11.

Proof. If $\rho \prec \sigma$, $\delta \prec \gamma$, and $\alpha \prec \beta$, then there exists a positive operator ρ' and $r > 0$ such that $\sigma = r\rho + \rho'$, a positive operator δ' and $d > 0$ such that $\gamma = d\delta + \delta'$ and a positive operator α' and $a > 0$ such that $\beta = a\alpha + \alpha'$ by Proposition 1. Then

$$\begin{aligned} f(\sigma \otimes \beta \otimes \gamma) &= (\epsilon^r \otimes 1 \otimes \epsilon^d)(\sigma \otimes \beta \otimes \gamma) \\ &= (\sigma \otimes 1 \otimes \gamma) \circ \beta \\ &= ((r\rho + \rho') \otimes 1 \otimes (d\delta + \delta')) \circ (a\alpha + \alpha') \\ &= (r\rho \otimes 1 \otimes d\delta) \circ (a\alpha + \alpha') + (\rho' \otimes 1 \otimes \delta') \circ (a\alpha + \alpha') \\ &= (r\rho \otimes 1 \otimes d\delta) \circ a\alpha + (r\rho \otimes 1 \otimes d\delta) \circ \alpha' + (\rho' \otimes 1 \otimes \delta') \circ (a\alpha + \alpha'), \\ f(\rho \otimes \alpha \otimes \delta) &= (\rho \otimes 1 \otimes \delta) \circ \alpha \end{aligned}$$

since $r, d, a \neq 0$, $\text{supp}(f(\rho \otimes \alpha \otimes \delta)) \subseteq \text{supp}(f(\sigma \otimes \beta \otimes \gamma))$, which by Proposition 1 proves the theorem.

6 Truth Theoretic Examples

We present several examples that demonstrate the application of the from-meanings-of-words-to-the-meaning-of-sentence map, where the initial meaning representations of words are density matrices, and explore how the hierarchy on nouns induced by their density matrix representations carry over to a hierarchy in the sentence space.

6.1 Entailment between nouns

Let “lions”, “sloths”, “plants” and “meat” have one dimensional representations in the noun space of our model:

$$\begin{aligned} \widehat{lions} &= |\overrightarrow{lions}\rangle\langle\overrightarrow{lions}| & \widehat{sloths} &= |\overrightarrow{sloths}\rangle\langle\overrightarrow{sloths}| \\ \widehat{meat} &= |\overrightarrow{meat}\rangle\langle\overrightarrow{meat}| & \widehat{plants} &= |\overrightarrow{plants}\rangle\langle\overrightarrow{plants}| \end{aligned}$$

Let the representation of “mammals” be a mixture of one dimensional representations of individual animals:

$$\widehat{mammals} = 1/2|\overrightarrow{lions}\rangle\langle\overrightarrow{lions}| + 1/2|\overrightarrow{sloths}\rangle\langle\overrightarrow{sloths}|$$

Notice that

$$\begin{aligned} N(\widehat{lions} || \widehat{mammals}) &= \text{tr}(\widehat{lions} \log \widehat{lions}) - \text{tr}(\widehat{lions} \log \widehat{mammals}) \\ &= \log 1 - \frac{1}{2} \log \frac{1}{2} = 1 \end{aligned}$$

Hence $R(\widehat{lions}, \widehat{mammals}) = 1/2$. For the other direction, since the intersection of the support of $\widehat{mammals}$ and the kernel of \widehat{lions} is non-empty, $R(\widehat{mammals}, \widehat{lions}) = 0$. This confirms that $\widehat{lions} \prec \widehat{mammals}$.

6.2 Entailment between sentences in one dimensional truth theoretic space

Consider a sentence space that is one dimensional, where 1 stands for true and 0 for false. Let sloths eat plants and lions eat meat; this is represented as follows

$$\begin{aligned} \widehat{eat} &= (|\overrightarrow{sloths}\rangle|\overrightarrow{plants}\rangle + |\overrightarrow{lions}\rangle|\overrightarrow{meat}\rangle)(\langle\overrightarrow{sloths}| \langle\overrightarrow{plants}| + \langle\overrightarrow{lions}| \langle\overrightarrow{meat}|) \\ &\approx (|\overrightarrow{sloths}\rangle\langle\overrightarrow{sloths}| \otimes |\overrightarrow{plants}\rangle\langle\overrightarrow{plants}|) + (|\overrightarrow{sloths}\rangle\langle\overrightarrow{lions}| \otimes |\overrightarrow{plants}\rangle\langle\overrightarrow{meat}|) + \\ &\quad (|\overrightarrow{lions}\rangle\langle\overrightarrow{sloths}| \otimes |\overrightarrow{meat}\rangle\langle\overrightarrow{plants}|) + (|\overrightarrow{lions}\rangle\langle\overrightarrow{lions}| \otimes |\overrightarrow{meat}\rangle\langle\overrightarrow{meat}|) \end{aligned}$$

The above is the density matrix representation of a pure composite state that relate “sloths” to “plants” and “lions” to “meat”. If we fix the bases $\{\overrightarrow{lions}, \overrightarrow{sloths}\}$ for N_1 , and $\{\overrightarrow{meat}, \overrightarrow{plants}\}$ for N_2 , we will have $\widehat{eat} : N_1 \otimes N_1 \rightarrow N_2 \otimes N_2$ with the following matrix representation:

$$\begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

“*Lions eat meat*” . This is a transitive sentence, so as before, it has the pregroup type: $nn^l sn^r n$. Explicit calculations for its meaning give:

$$\begin{aligned} &(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{lions} \otimes \widehat{eat} \otimes \widehat{meat}) \\ &= \langle\overrightarrow{lions}|\overrightarrow{sloths}\rangle^2 \langle\overrightarrow{plants}|\overrightarrow{meat}\rangle^2 + \\ &\quad \langle\overrightarrow{lions}|\overrightarrow{sloths}\rangle \langle\overrightarrow{lions}|\overrightarrow{lions}\rangle \langle\overrightarrow{meat}|\overrightarrow{meat}\rangle \langle\overrightarrow{plants}|\overrightarrow{meat}\rangle + \\ &\quad \langle\overrightarrow{lions}|\overrightarrow{lions}\rangle \langle\overrightarrow{lions}|\overrightarrow{sloths}\rangle \langle\overrightarrow{meat}|\overrightarrow{meat}\rangle \langle\overrightarrow{plants}|\overrightarrow{meat}\rangle + \\ &\quad \langle\overrightarrow{lions}|\overrightarrow{lions}\rangle^2 \langle\overrightarrow{meat}|\overrightarrow{meat}\rangle^2 \\ &= 1 \end{aligned}$$

“*Sloths eat meat*” . This sentence has a very similar calculation to the one above with the resulting meaning:

$$(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{sloths} \otimes \widehat{eat} \otimes \widehat{meat}) = 0$$

“*Mammals eat meat*” . This sentence has the following meaning calculation:

$$\begin{aligned} &(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{mammals} \otimes \widehat{eat} \otimes \widehat{meat}) = \\ &(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)((\frac{1}{2}\widehat{lions} + \frac{1}{2}\widehat{sloths}) \otimes \widehat{eat} \otimes \widehat{meat}) = \\ &\frac{1}{2}(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{lions} \otimes \widehat{eat} \otimes \widehat{meat}) + \frac{1}{2}(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{sloths} \otimes \widehat{eat} \otimes \widehat{meat}) = \frac{1}{2} \end{aligned}$$

The resulting meaning of this sentence is a mixture of “lions eat meat”, which is true, and “sloths eat meat” which is false. Thus the value 1/2 can be interpreted as being neither completely true or completely false: the sentence “mammals eat meat” is true for certain mammals and false for others.

6.3 Entailment between sentences in two dimensional truth theoretic space

The two dimensional truth theoretic space is set as follows:

$$true \equiv |0\rangle \equiv \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad false \equiv |1\rangle \equiv \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

The corresponding *true* and *false* density matrices are $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$.

In the two dimensional space, the representation of “eats” is set as follows. Let $A = \{lions, sloths\}$ and $B = \{meat, plants\}$, then

$$\widehat{eat} \equiv \sum_{\substack{a_1, a_2 \in A \\ b_1, b_2 \in B}} |\vec{a}_1\rangle\langle \vec{a}_2| \otimes |\vec{x}\rangle\langle \vec{x}| \otimes |\vec{b}_1\rangle\langle \vec{b}_2|$$

where

$$|x\rangle \equiv \begin{cases} |0\rangle & \text{if } |a_1\rangle|b_1\rangle, |a_2\rangle|b_2\rangle \in \{|\overrightarrow{lions}\rangle|\overrightarrow{meat}\rangle, |\overrightarrow{sloths}\rangle|\overrightarrow{plants}\rangle\} \\ |1\rangle & \text{otherwise} \end{cases}$$

The generalized matrix representation of this verb in the spirit of [14] is:

$$\begin{pmatrix} 1 & 0 & 0 & 1 & | & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & | & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & | & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & | & 0 & 1 & 1 & 0 \end{pmatrix}$$

“*Lions eat meat*”. The calculation for the meaning of this sentence is almost exactly the same as the case of the one dimensional meaning, only the result is not the scalar that stands for *true* but its density matrix:

$$(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{lions} \otimes \widehat{eat} \otimes \widehat{meat}) = |0\rangle\langle 0|$$

“*Sloths eat meat*”. Likewise, the calculation for the meaning of this sentence returns *false*:

$$(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{sloths} \otimes \widehat{eat} \otimes \widehat{meat}) = |1\rangle\langle 1|$$

“*Mammals eat meat*”. As we saw before, this sentence has the meaning that is the mixture of “Lions eat meat” and “Sloths eat meat”; here, this is expressed as follows:

$$\begin{aligned} & (\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{mammals} \otimes \widehat{eat} \otimes \widehat{meat}) \\ &= \frac{1}{2}(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{lions} \otimes \widehat{eat} \otimes \widehat{meat}) + \frac{1}{2}(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{sloths} \otimes \widehat{eat} \otimes \widehat{meat}) \\ &= \frac{1}{2}|1\rangle\langle 1| + \frac{1}{2}|0\rangle\langle 0| \end{aligned}$$

So in a two dimensional truth theoretic model, “Mammals eat meat” give the completely mixed state in the sentence space, which has maximal entropy. This is equivalent to

saying that we have no real knowledge whether mammals in general eat meat or not. Even if we are completely certain about whether individual mammals that span our space for “mammals” eat meat, this information differs uniformly within the members of the class, so we cannot generalize.

Already with a two dimensional truth theoretic model, the relation $\widehat{lions} \prec \widehat{mammals}$ carries over to sentences. To see this, first note that we have

$$\begin{aligned} N(\widehat{lions\ eat\ meat} | \widehat{mammals\ eat\ meat}) &= N\left(|0\rangle\langle 0| \left\| \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| \right\| \right) \\ &= (|0\rangle\langle 0|) \log(|0\rangle\langle 0|) - (|0\rangle\langle 0|) \log\left(\frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|\right) = 1 \end{aligned}$$

In the other direction, we have $N(\widehat{mammals\ eat\ meat} | \widehat{lions\ eat\ meat}) = \infty$, since the intersection of the support of the first argument and the kernel of the second argument is non-trivial. These lead to the following representativeness results between sentences:

$$\begin{aligned} R(\widehat{lions\ eat\ meat}, \widehat{mammals\ eat\ meat}) &= 1/2 \\ R(\widehat{mammals\ eat\ meat} | \widehat{lions\ eat\ meat}) &= 0 \end{aligned}$$

As a result we obtain:

$$\widehat{lions\ eat\ meat} \prec \widehat{mammals\ eat\ meat}$$

Since these two sentences share the same verb phrase, from-meaning-of-words-to-the-meaning-of-sentence map carries the hyponymy relation in the subject words of the respective sentences to the resulting sentence meanings. By using the density matrix representations of word meanings together with the categorical map from the meanings of words to the meanings of sentences, the knowledge that a lion is an animal lets us infer that “mammals eat meat” implies “lions eat meat”:

$$(\widehat{lions} \prec \widehat{mammals}) \rightarrow (\widehat{lions\ eat\ meat} \prec \widehat{mammals\ eat\ meat})$$

“Dogs eat meat”. To see how the completely mixed state differs from a perfectly correlated but pure state in the context of linguistic meaning, consider a new noun $\widehat{dog} = |\widehat{dog}\rangle\langle\widehat{dog}|$ and redefine eat in terms of the bases $\{\widehat{lions}, \widehat{dogs}\}$ and $\{\widehat{meat}, \widehat{plants}\}$, so that it will reflect the fact that dogs eat *both* meat and plants. We define “eat” so that it results in the value of being “half-true half-false” when it takes “dogs” as subject and “meat” or “plants” as object. The value “half-true half-false” is the superposition of *true* and *false*: $\frac{1}{2}|0\rangle + \frac{1}{2}|1\rangle$. With this assumptions, \widehat{eat} will still be a pure state with the following representation in **FVect**:

$$\begin{aligned} |\widehat{eat}\rangle &= |\widehat{lions}\rangle \otimes |0\rangle \otimes |\widehat{meat}\rangle + |\widehat{lions}\rangle \otimes |1\rangle \otimes |\widehat{plants}\rangle + \\ &|\widehat{dogs}\rangle \otimes \left(\frac{1}{2}|0\rangle + \frac{1}{2}|1\rangle\right) \otimes |\widehat{meat}\rangle + |\widehat{dogs}\rangle \otimes \left(\frac{1}{2}|0\rangle + \frac{1}{2}|1\rangle\right) \otimes |\widehat{plants}\rangle \end{aligned}$$

Hence, the density matrix representation of “eat” becomes:

$$\widehat{eat} = |\overrightarrow{eat}\rangle\langle\overrightarrow{eat}|$$

The calculation for the meaning of the sentence is as follows:

$$\begin{aligned} & (\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{dogs} \otimes \widehat{eat} \otimes \widehat{meat}) \\ &= (\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(|\overrightarrow{dogs}\rangle\langle\overrightarrow{dogs}| \otimes |\overrightarrow{eat}\rangle\langle\overrightarrow{eat}| \otimes |\overrightarrow{meat}\rangle\langle\overrightarrow{meat}|) \\ &= \left(\frac{1}{2}|0\rangle + \frac{1}{2}|1\rangle\right)\left(\frac{1}{2}\langle 0| + \frac{1}{2}\langle 1|\right) \end{aligned}$$

So in this case, we are certain that it is half-true and half-false that dogs eat meat. This is in contrast with the completely mixed state we got from “Mammals eat meat”, for which the truth or falsity of the sentence was entirely unknown.

“Mammals eat meat”, again. Let “mammals” now be defined as:

$$\widehat{mammals} = \frac{1}{2}\widehat{lions} + \frac{1}{2}\widehat{dogs}$$

The calculation for the meaning of this sentence gives:

$$\begin{aligned} & (\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{mammals} \otimes \widehat{eat} \otimes \widehat{meat}) \\ &= \frac{1}{2}(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{lions} \otimes \widehat{eat} \otimes \widehat{meat}) + \frac{1}{2}(\epsilon_N^l \otimes 1_S \otimes \epsilon_N^r)(\widehat{dogs} \otimes \widehat{eat} \otimes \widehat{meat}) \\ &= \frac{3}{4}|0\rangle\langle 0| + \frac{1}{4}|0\rangle\langle 1| + \frac{1}{4}|1\rangle\langle 0| + \frac{1}{4}|1\rangle\langle 1| \end{aligned}$$

This time the resulting sentence representation is not completely mixed. This means that we can generalize the knowledge we have from the specific instances of mammals to the entire class to some extent, but still we cannot generalize completely. This is a mixed state, which indicates that even if the sentence is closer to *true* than to *false*, the degree of truth isn’t homogeneous throughout the elements of the class. The non-zero non-diagonals indicate that it is also partially correlated, which means that there are some instances of “mammals” for which this sentence is true to a degree, but not completely. The relative similarity measures between *true* and *false* and the sentence can be calculated explicitly using fidelity:

$$F(|1\rangle\langle 1|, \widehat{mammals} \widehat{eat} \widehat{meat}) = \langle 1|\widehat{mammals} \widehat{eat} \widehat{meat}|1\rangle = \frac{1}{4}$$

$$F(|0\rangle\langle 0|, \widehat{mammals} \widehat{eat} \widehat{meat}) = \langle 0|\widehat{mammals} \widehat{eat} \widehat{meat}|0\rangle = \frac{3}{4}$$

Notice that these values are different from the values for the representativeness for truth and falsity of the sentence, even though they are proportional: the more

representative their density matrices, the more similar the sentences are to each other. For example, we have:

$$N(|1\rangle\langle 1| \parallel \widehat{\text{mammals eat meat}}) = \text{tr}(|1\rangle\langle 1| \log(|1\rangle\langle 1|)) - \text{tr}(|1\rangle\langle 1| \log(\frac{3}{4}|0\rangle\langle 0| + \frac{1}{4}|0\rangle\langle 1| + \frac{1}{4}|1\rangle\langle 0| + \frac{1}{4}|1\rangle\langle 1|)) \approx 2$$

Hence, $R(|1\rangle\langle 1| \parallel \widehat{\text{mammals eat meat}}) \approx .33$. On the other hand:

$$N(|0\rangle\langle 0| \parallel \widehat{\text{mammals eat meat}}) = \text{tr}(|0\rangle\langle 0| \log(|0\rangle\langle 0|)) - \text{tr}(|0\rangle\langle 0| \log(\frac{3}{4}|0\rangle\langle 0| + \frac{1}{4}|0\rangle\langle 1| + \frac{1}{4}|1\rangle\langle 0| + \frac{1}{4}|1\rangle\langle 1|)) \approx 0.41$$

Hence, $R(|0\rangle\langle 0| \parallel \widehat{\text{mammals eat meat}}) \approx 0.71$

7 Distributional Examples

The goal of this section is to show how one can obtain density matrices for words using lexical taxonomies and co-occurrence frequencies counted from corpora of text. We show how these density matrices are used in example sentences and how the density matrices of their meanings look like. We compute the representativeness formula for these sentences to provide a proof of concept that this measure does makes sense for data harvested from corpora distributionally and that its application is not restricted to truth-theoretic models. Implementing these constructions on real data and validating them on large scale datasets constitute work in progress.

7.1 Entailment between nouns

Suppose we have a noun space N . Let the subspace relevant for this part of the example be spanned by lemmas *pub*, *pitcher*, *tonic*. Assume that the (non-normalized version of the) vectors of the atomic words *lager* and *ale* in this subspace are as follows:

$$\begin{aligned} \overrightarrow{\text{lager}} &= 6 \times \overrightarrow{\text{pub}} + 5 \times \overrightarrow{\text{pitcher}} + 0 \times \overrightarrow{\text{tonic}} \\ \overrightarrow{\text{ale}} &= 7 \times \overrightarrow{\text{pub}} + 3 \times \overrightarrow{\text{pitcher}} + 0 \times \overrightarrow{\text{tonic}} \end{aligned}$$

Suppose further that we are given taxonomies such as ‘beer = lager + ale’, harvested from a resource such as WordNet. Atomic words (i.e. leafs of the taxonomy), correspond to *pure* states and their density matrices are the projections onto the one dimensional subspace spanned by $|\overrightarrow{w}\rangle\langle\overrightarrow{w}|$. Non-atomic words (such as *beer*) are also density matrices, harvested from the corpus using a feature-based method similar to that of [12]. This is done by counting (and normalising) the frequency of times a word has co-occurred with a subset B of bases in a window in which other bases (the ones not in B) have not occurred.

Formally, for a subset of bases $\{b_1, b_2, \dots, b_n\}$, we collect co-ordinates C_{ij} for each tuple $|b_i\rangle|b_j\rangle$ and build the density matrix $\sum_{ij} C_{ij} |b_i\rangle|b_j\rangle$.

For example, suppose we see *beer* six times with just *pub*, seven times with both *pub* and *pitcher*, and none-whatsoever with *tonic*. Its corresponding density matrix will be as follows:

$$\begin{aligned} \widehat{beer} &= 6 \times |\overrightarrow{pub}\rangle\langle\overrightarrow{pub}| + 7 \times (|\overrightarrow{pub}\rangle + |\overrightarrow{pitcher}\rangle)(\langle\overrightarrow{pub}| + \langle\overrightarrow{pitcher}|) \\ &= 13 \times |\overrightarrow{pub}\rangle\langle\overrightarrow{pub}| + 7 \times |\overrightarrow{pub}\rangle\langle\overrightarrow{pitcher}| + 7 \times |\overrightarrow{pitcher}\rangle\langle\overrightarrow{pub}| + 7 \times |\overrightarrow{pitcher}\rangle\langle\overrightarrow{pitcher}| \end{aligned}$$

To calculate the similarity and representativeness of the word pairs, we first normalize them via the operation $\frac{\rho}{\text{Tr}\rho}$, then apply the corresponding formulae. For example, the degree of similarity between ‘beer’ and ‘lager’ using fidelity is as follows:

$$\text{Tr}\sqrt{\widehat{lager}^{\frac{1}{2}} \cdot \widehat{beer} \cdot \widehat{lager}^{\frac{1}{2}}} = 0.93$$

The degree of entailment $lager \prec beer$ is 0.82 as computed as follows:

$$\frac{1}{1 + \text{Tr}(\widehat{lager} \cdot \log(\widehat{lager}) - \widehat{lager} \cdot \log(\widehat{beer}))} = 0.82$$

The degree of entailment $beer \prec lager$ is 0, like one would expect.

7.2 Entailment between sentences

To see how the entailment between sentences follows from the entailment between words, consider example sentences ‘*Psychiatrist is drinking lager*’ and ‘*Doctor is drinking beer*’. For the sake of brevity, we assume the meanings of *psychiatrist* and *doctor* are mixtures of basis elements, as follows:

$$\begin{aligned} \widehat{psychiatrist} &= 2 \times |\overrightarrow{patient}\rangle\langle\overrightarrow{patient}| + 5 \times |\overrightarrow{mental}\rangle\langle\overrightarrow{mental}| \\ \widehat{doctor} &= 5 \times |\overrightarrow{patient}\rangle\langle\overrightarrow{patient}| + 2 \times |\overrightarrow{mental}\rangle\langle\overrightarrow{mental}| + 3 \times |\overrightarrow{surgery}\rangle\langle\overrightarrow{surgery}| \end{aligned}$$

The similarity between *psychiatrist* and *doctor* is:

$$S(\widehat{psychiatrist}, \widehat{doctor}) = S(\widehat{doctor}, \widehat{psychiatrist}) = 0.76$$

The representativeness between them is:

$$R(\widehat{psychiatrist}, \widehat{doctor}) = 0.49 \quad R(\widehat{doctor}, \widehat{psychiatrist}) = 0$$

We build matrices for the verb *drink* following the method of [15]. Intuitively this is as follows: the value in entry (i, j) of this matrix will reflect how typical it is for the verb to have a subject related to the i th basis and an object related to the j th basis. We assume that the small part of the matrix that interests us for this example is as follows:

<i>drink</i>	<i>pub</i>	<i>pitcher</i>	<i>tonic</i>
<i>patient</i>	4	5	3
<i>mental</i>	6	3	2
<i>surgery</i>	1	2	1

This representation can be seen as a pure state living in a second order tensor. Therefore the density matrix representation of the same object is $\widehat{drink} = |\widehat{drink}\rangle\langle\widehat{drink}|$, a fourth order tensor. Lifting the simplifications introduced in [15] from vectors to density matrices, we obtain the following linear algebraic closed forms for the meaning of the sentences:

$$\begin{aligned} \widehat{Psychiatrist\ is\ drinking\ lager} &= \widehat{drink} \odot (\widehat{psychiatrist} \otimes \widehat{lager}) \\ \widehat{Doctor\ is\ drinking\ beer} &= \widehat{drink} \odot (\widehat{doctor} \otimes \widehat{beer}) \end{aligned}$$

Applying the fidelity and representativeness formulae to sentence representations, we obtain the following values:

$$\begin{aligned} S(\widehat{Psychiatrist\ is\ drinking\ lager}, \widehat{Doctor\ is\ drinking\ beer}) &= 0.81 \\ R(\widehat{Psychiatrist\ is\ drinking\ lager}, \widehat{Doctor\ is\ drinking\ beer}) &= 0.53 \\ R(\widehat{Doctor\ is\ drinking\ beer}, \widehat{Psychiatrist\ is\ drinking\ lager}) &= 0 \end{aligned}$$

From the relations $\widehat{psychiatrist} \prec \widehat{doctor}$ and $\widehat{lager} \prec \widehat{beer}$ we obtain the desired entailment between sentences:

$$\widehat{Psychiatrist\ is\ drinking\ lager} \prec \widehat{Doctor\ is\ drinking\ beer}.$$

The entailment between these two sentences follows from the entailment between their subjects and the entailment between their objects. In the examples that we have considered so far, the verbs of sentences are the same. This is not a necessity. One can have entailment between sentences that do not have the same verbs, but where the verbs entail each other, examples can be found in [2]. The reason we do not present such cases here is lack of space.

8 Conclusion and Future Work

The often stated long term goal of compositional distributional models is to merge distributional and formal semantics. However, what formal and distributional semantics *do* with the resulting meaning representations is quite different. Distributional semanticists care about *similarity* while formal semanticists aim to capture *truth* and *inference*. In this work we presented a theory of meaning using basic objects that will not confine us to the realm of only distributional or only formal semantics. The immediate next step is to develop methods for obtaining density matrix representations of words from corpus, that are more robust to statistical noise, and testing the usefulness of the theory in large scale experiments.

The problem of integrating function words such as ‘and’, ‘or’, ‘not’, ‘every’ into a distributional setting has been notoriously hard. We hope that the characterization of compositional distributional entailment on these very simple types of sentences will provide a foundation on which we can define representations of these function words, and develop a more logical theory of compositional distributional meaning.

Bibliography

- [1] S. Abramsky and B. Coecke. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science*, pages 415–425. IEEE Computer Science Press, 2004. arXiv:quant-ph/0402130.
- [2] Esma Balkır. Using density matrices in a compositional distributional model of meaning. Master’s thesis, University of Oxford, 2014.
- [3] Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. Entailment above the word level in distributional semantics. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 23–32. Association for Computational Linguistics, 2012.
- [4] Islam Beltagy, Cuong Chau, Gemma Boleda, Dan Garrette, Katrin Erk, and Raymond Mooney. Montague meets markov: Deep semantics with probabilistic logical form. In *Second Joint Conference on Lexical and Computational Semantics*, volume 1, pages 11–21. Association of Computational Linguists, 2013.
- [5] William Blacoe, Elham Kashefi, and Mirella Lapata. A quantum-theoretic approach to distributional semantics. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 847–857, 2013.
- [6] Peter D Bruza and Richard Cole. Quantum logic of semantic space: An exploratory investigation of context effects in practical reasoning. *We Will Show Them: Essays in Honour of Dov Gabbay*, pages 339–361, 2005.
- [7] Stephen Clark, Bob Coecke, and Mehrnoosh Sadrzadeh. A compositional distributional model of meaning. *Proceedings of the Second Symposium on Quantum Interaction (QI-2008)*, pages 133–140, 2008.
- [8] Daoud Clarke. Context-theoretic semantics for natural language: an overview. In *Proceedings of the workshop on geometrical models of natural language semantics*, pages 112–119. Association for Computational Linguistics, 2009.
- [9] Bob Coecke. Quantum picturalism. *Contemporary physics*, 51(1):59–83, 2010.
- [10] Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis*, 36, 2010.
- [11] John R. Firth. A Synopsis of Linguistic Theory, 1930-1955. *Studies in Linguistic Analysis*, pages 1–32, 1957.
- [12] Maayan Geffet and Ido Dagan. The distributional inclusion hypotheses and lexical entailment. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 107–114. Association for Computational Linguistics, 2005.
- [13] Oren Glickman, Ido Dagan, and Moshe Koppel. Web based probabilistic textual entailment. In *Proceedings of the PASCAL Challenges Workshop on Recognizing Textual Entailment*, 2005.
- [14] Edward Grefenstette. *Category-Theoretic Quantitative Compositional Distributional Models of Natural Language Semantics*. PhD thesis, University of Oxford, 2013.

- [15] Edward Grefenstette and Mehrnoosh Sadrzadeh. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1394–1404. Association for Computational Linguistics, 2011.
- [16] Aurélie Herbelot and Mohan Ganesalingam. Measuring semantic content in distributional vectors. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 440–445. Association for Computational Linguistics, 2013.
- [17] Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. Directional distributional similarity for lexical inference. *Natural Language Engineering*, 16(04):359–389, 2010.
- [18] Joachim Lambek. Type grammars as pregroups. *Grammars*, 4(1):21–39, 2001.
- [19] Alessandro Lenci and Giulia Benotto. Identifying hypernyms in distributional semantic spaces. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, volume 2, pages 75–79. Association for Computational Linguistics, 2012.
- [20] Dekang Lin. An information-theoretic definition of similarity. In *Proceedings of the International Conference on Machine Learning*, pages 296–304, 1998.
- [21] Bill MacCartney and Christopher D. Manning. Natural logic for textual inference. In *ACL Workshop on Textual Entailment and Paraphrasing*. Association for Computational Linguistics, 2007.
- [22] Robin Piedeleu, Dimitri Kartsaklis, Bob Coecke, and Mehrnoosh Sadrzadeh. Open system categorical quantum semantics in natural language processing. *arXiv preprint arXiv:1502.00831*, 2015.
- [23] Enrico Santus, Alessandro Lenci, Qin Lu, and Sabine Schulte Im Walde. Chasing hypernyms in vector spaces with entropy. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, volume 2, pages 38–42, 2014.
- [24] Peter Selinger. Dagger compact closed categories and completely positive maps. *Electronic Notes in Theoretical Computer Science*, 170:139–163, 2007.
- [25] Alessandro Sordoni, Jian-Yun Nie, and Yoshua Bengio. Modeling term dependencies with quantum language models for ir. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 653–662. Association for Computational Linguistics, 2013.
- [26] Cornelis Joost Van Rijsbergen. *The geometry of information retrieval*. Cambridge University Press, 2004.
- [27] Julie Weeds, David Weir, and Diana McCarthy. Characterising measures of lexical distributional similarity. In *Proceedings of the 20th international conference on Computational Linguistics*, number 1015. Association for Computational Linguistics, 2004.