# Towards a Formal Distributional Semantics: Simulating Logical Calculi with Tensors 

Edward Grefenstette<br>University of Oxford<br>Department of Computer Science<br>Wolfson Building, Parks Road<br>Oxford OX1 3QD, UK<br>edward.grefenstette@cs.ox.ac.uk


#### Abstract

The development of compositional distributional models of semantics reconciling the empirical aspects of distributional semantics with the compositional aspects of formal semantics is a popular topic in the contemporary literature. This paper seeks to bring this reconciliation one step further by showing how the mathematical constructs commonly used in compositional distributional models, such as tensors and matrices, can be used to simulate different aspects of predicate logic. This paper discusses how the canonical isomorphism between tensors and multilinear maps can be exploited to simulate a full-blown quantifier-free predicate calculus using tensors. It provides tensor interpretations of the set of logical connectives required to model propositional calculi. It suggests a variant of these tensor calculi capable of modelling quantifiers, using few non-linear operations. It finally discusses the relation between these variants, and how this relation should constitute the subject of future work.


## 1 Introduction

The topic of compositional distributional semantics has been growing in popularity over the past few years. This emerging sub-field of natural language semantic modelling seeks to combine two seemingly orthogonal approaches to modelling the meaning of words and sentences, namely formal semantics and distributional semantics.

These approaches, summarised in Section 2, differ in that formal semantics, on the one hand, pro-
vides a neatly compositional picture of natural language meaning, reducing sentences to logical representations; one the other hand, distributional semantics accounts for the ever-present ambiguity and polysemy of words of natural language, and provides tractable ways of learning and comparing word meanings based on corpus data.
Recent efforts, some of which are briefly reported below, have been made to unify both of these approaches to language modelling to produce compositional distributional models of semantics, leveraging the learning mechanisms of distributional semantics, and providing syntax-sensitive operations for the production of representations of sentence meaning obtained through combination of corpus-inferred word meanings. These efforts have been met with some success in evaluations such as phrase similarity tasks (Mitchell and Lapata, 2008; Mitchell and Lapata, 2009; Grefenstette and Sadrzadeh, 2011; Kartsaklis et al., 2012), sentiment prediction (Socher et al., 2012), and paraphrase detection (Blacoe and Lapata, 2012).
While these developments are promising with regard to the goal of obtaining learnable-yetstructured sentence-level representations of language meaning, part of the motivation for unifying formal and distributional models of semantics has been lost. The compositional aspects of formal semantics are combined with the corpus-based empirical aspects of distributional semantics in such models, yet the logical aspects are not. But it is these logical aspects which are so appealing in formal semantic models, and therefore it would be desirable to replicate the inferential powers of logic within
compositional distributional models of semantics.
In this paper, I make steps towards addressing this lost connection with logic in compositional distributional semantics. In Section 2, I provide a brief overview of formal and distributional semantic models of meaning. In Section 3, I give mathematical foundations for the rest of the paper by introducing tensors and tensor contraction as a way of modelling multilinear functions. In Section 4, I discuss how predicates, relations, and logical atoms of a quantifier-free predicate calculus can be modelled with tensors. In Section 5, I present tensorial representations of logical operations for a complete propositional calculus. In Section 6, I discuss a variant of the predicate calculus from Section 4 aimed at modelling quantifiers within such tensorbased logics, and the limits of compositional formalisms based only on multilinear maps. I conclude, in Section 7, by suggesting directions for further work based on the contents of this paper.

This paper does not seek to address the question of how to determine how words should be translated into predicates and relations in the first place, but rather shows how such predicates and relations can be modelled using multilinear algebra. As such, it can be seen as a general theoretical contribution which is independent from the approaches to compositional distributional semantics it can be applied to. It is directly compatible with the efforts of Coecke et al. (2010) and Grefenstette et al. (2013), discussed below, but is also relevant to any other approach making use of tensors or matrices to encode semantic relations.

## 2 Related work

Formal semantics, from the Montagovian school of thought (Montague, 1974; Dowty et al., 1981), treats natural languages as programming languages which compile down to some formal language such as a predicate calculus. The syntax of natural languages, in the form of a grammar, is augmented by semantic interpretations, in the form of expressions from a higher order logic such as the lambda-beta calculus. The parse of a sentence then determines the combinations of lambda-expressions, the reduction of which yields a well-formed formula of a predicate calculus, corresponding to the semantic repre-
sentation of the sentence. A simple formal semantic model is illustrated in Figure 1.

| Syntactic Analysis | Semantic Interpretation |
| :--- | :--- |
| $\mathrm{S} \Rightarrow \mathrm{NP}$ VP | $\llbracket V P \rrbracket(\llbracket N P \rrbracket)$ |
| $\mathrm{NP} \Rightarrow$ cats, milk, etc. | $\llbracket \mathrm{cats} \rrbracket, \llbracket \mathrm{milk} \rrbracket, \ldots$ |
| $\mathrm{VP} \Rightarrow \mathrm{Vt}$ NP | $\llbracket V t \rrbracket(\llbracket N P \rrbracket)$ |
| $\mathrm{Vt} \Rightarrow$ like, hug, etc. | $\lambda y x . \llbracket$ like $\rrbracket(x, y), \ldots$ |



Figure 1: A simple formal semantic model.
Formal semantic models are incredibly powerful, in that the resulting logical representations of sentences can be fed to automated theorem provers to perform textual inference, consistency verification, question answering, and a host of other tasks which are well developed in the literature (e.g. see (Loveland, 1978) and (Fitting, 1996)). However, the sophistication of such formal semantic models comes at a cost: the complex set of rules allowing for the logical interpretation of text must either be provided a priori, or learned. Learning such representations is a complex task, the difficulty of which is compounded by issues of ambiguity and polysemy which are pervasive in natural languages.

In contrast, distributional semantic models, best summarised by the dictum of Firth (1957) that "You shall know a word by the company it keeps," provide an elegant and tractable way of learning semantic representations of words from text. Word meanings are modelled as high-dimensional vectors in large semantic vector spaces, the basis elements of which correspond to contextual features such as other words from a lexicon. Semantic vectors for words are built by counting how many time a target word occurs within a context (e.g. within $k$ words of select words from the lexicon). These context counts are then normalised by a term frequencyinverse document frequency-like measure (e.g. TFIDF, pointwise mutual information, ratio of probabilities), and are set as the basis weights of the vector representation of the word's meaning. Word vectors can then be compared using geometric distance


Figure 2: A simple distributional semantic model.
metrics such as cosine similarity, allowing us to determine the similarity of words, cluster semantically related words, and so on. Excellent overviews of distributional semantic models are provided by Curran (2004) and Mitchell (2011). A simple distributional semantic model showing the spacial representation of words 'dog', 'cat' and 'snake' within the context of feature words 'pet', 'furry', and 'stroke' is shown in Figure 2.

Distributional semantic models have been successfully applied to tasks such as word-sense discrimination (Schütze, 1998), thesaurus extraction (Grefenstette, 1994), and automated essay marking (Landauer and Dumais, 1997). However, while such models provide tractable ways of learning and comparing word meanings, they do not naturally scale beyond word length. As recently pointed out by Turney (2012), treating larger segments of texts as lexical units and learning their representations distributionally (the 'holistic approach') violates the principle of linguistic creativity, according to which we can formulate and understand phrases which we've never observed before, provided we know the meaning of their parts and how they are combined. As such, distributional semantics makes no effort to account for the compositional nature of language like formal semantics does, and ignores issues relating to syntactic and relational aspects of language.

Several proposals have been put forth over the last few years to provide vector composition functions for distributional models in order to introduce compositionality, thereby replicating some of the as-
pects of formal semantics while preserving learnability. Simple operations such as vector addition and multiplication, with or without scalar or matrix weights (to take word order or basic relational aspects into account), have been suggested (Zanzotto et al., 2010; Mitchell and Lapata, 2008; Mitchell and Lapata, 2009).

Smolensky (1990) suggests using the tensor product of word vectors to produce representations that grow with sentence complexity. Clark and Pulman (2006) extend this approach by including basis vectors standing for dependency relations into tensor product-based representations. Both of these tensor product-based approaches run into dimensionality problems as representations of sentence meaning for sentences of different lengths or grammatical structure do not live in the same space, and thus cannot directly be compared. Coecke et al. (2010) develop a framework using category theory, solving this dimensionality problem of tensor-based models by projecting tensored vectors for sentences into a unique vector space for sentences, using functions dynamically generated by the syntactic structure of the sentences. In presenting their framework, which partly inspired this paper, they describe how a verb can be treated as a logical relation using tensors in order to evaluate the truth value of a simple sentence, as well as how negation can be modelled using matrices.

A related approach, by Baroni and Zamparelli (2010), represents unary relations such as adjectives as matrices learned by linear regression from corpus data, and models adjective-noun composition as matrix-vector multiplication. Grefenstette et al. (2013) generalise this approach to relations of any arity and relate it to the framework of Coecke et al. (2010) using a tensor-based approach to formal semantic modelling similar to that presented in this paper.

Finally, Socher et al. (2012) apply deep learning techniques to model syntax-sensitive vector composition using non-linear operations, effectively turning parse trees into multi-stage neural networks. Socher shows that the non-linear activation function used in such a neural network can be tailored to replicate the behaviour of basic logical connectives such as conjunction and negation.

## 3 Tensors and multilinear maps

Tensors are the mathematical objects dealt with in multilinear algebra just as vectors and matrices are the objects dealt with in linear algebra. In fact, tensors can be seen as generalisations of vectors and matrices by introducing the notion of tensor rank. Let the rank of a tensor be the number of indices required to describe a vector/matrix-like object in sum notation. A vector $\mathbf{v}$ in a space $V$ with basis $\left\{\mathbf{b}_{i}^{V}\right\}_{i}$ can be written as the weighted sum of the basis vectors:

$$
\mathbf{v}=\sum_{i} c_{i}^{v} \mathbf{b}_{i}^{V}
$$

where the $c_{i}^{v}$ elements are the scalar basis weights of the vector. Being fully described with one index, vectors are rank 1 tensors. Similarly, a matrix $\mathbf{M}$ is an element of a space $V \otimes W$ with basis $\left\{\left(\mathbf{b}_{i}^{V}, \mathbf{b}_{j}^{W}\right)\right\}_{i j}$ (such pairs of basis vectors of $V$ and $W$ are commonly written as $\left\{\mathbf{b}_{i}^{V} \otimes \mathbf{b}_{j}^{W}\right\}_{i j}$ in multilinear algebra). Such matrices are rank 2 tensors, as they can be fully described using two indices (one for rows, one for columns):

$$
\mathbf{M}=\sum_{i j} c_{i j}^{M} \mathbf{b}_{i}^{V} \otimes \mathbf{b}_{j}^{W}
$$

where the scalar weights $c_{i j}^{M}$ are just the $i j$ th elements of the matrix.

A tensor $\mathbf{T}$ of rank $k$ is just a geometric object with a higher rank. Let $T$ be a member of $V_{1} \otimes \ldots \otimes V_{k}$; we can express $T$ as follows, using $k$ indices $\alpha_{1} \ldots \alpha_{k}$ :

$$
\mathbf{T}=\sum_{\alpha_{1} \ldots \alpha_{k}} c_{\alpha_{1} \ldots \alpha_{k}}^{T} \mathbf{b}_{\alpha_{1}}^{V_{1}} \otimes \ldots \otimes \mathbf{b}_{\alpha_{k}}^{V_{k}}
$$

In this paper, we will be dealing with tensors of rank 1 (vectors), rank 2 (matrices) and rank 3, which can be pictured as cuboids (or a matrix of matrices).

Tensor contraction is an operation which allows us to take two tensors and produce a third. It is a generalisation of inner products and matrix multiplication to tensors of higher ranks. Let $\mathbf{T}$ be a tensor in $V_{1} \otimes \ldots \otimes V_{j} \otimes V_{k}$ and $\mathbf{U}$ be a tensor in $V_{k} \otimes V_{m} \otimes \ldots \otimes V_{n}$. The contraction of these tensors, written $\mathbf{T} \times \mathbf{U}$, corresponds to the following calculation:

$$
\begin{aligned}
& \mathbf{T} \times \mathbf{U}= \\
& \quad \sum_{\alpha_{1} \ldots \alpha_{n}} c_{\alpha_{1} \ldots \alpha_{k}}^{T} c_{\alpha_{k} \ldots \alpha_{n}}^{U} \mathbf{b}_{\alpha_{1}}^{V_{1}} \otimes \ldots \otimes \mathbf{b}_{\alpha_{j}}^{V_{j}} \otimes \mathbf{b}_{\alpha_{m}}^{V_{m}} \otimes \ldots \otimes \mathbf{b}_{\alpha_{n}}^{V_{n}}
\end{aligned}
$$

Tensor contraction takes a tensor of rank $k$ and a tensor of rank $n-k+1$ and produces a tensor of
rank $n-1$, corresponding to the sum of the ranks of the input tensors minus 2. The tensors must satisfy the following restriction: the left tensor must have a rightmost index spanning the same number of dimensions as the leftmost index of the right tensor. This is similar to the restriction that a $m$ by $n$ matrix can only be multiplied with a $p$ by $q$ matrix if $n=p$, i.e. if the index spanning the columns of the first matrix covers the same number of columns as the index spanning the rows of the second matrix covers rows. Similarly to how the columns of one matrix 'merge' with the rows of another to produce a third matrix, the part of the first tensor spanned by the index $k$ merges with the part of the second tensor spanned by $k$ by 'summing through' the shared basis elements $\mathbf{b}_{\alpha_{k}}^{V_{k}}$ of each tensor. Each tensor therefore loses a rank while being joined, explaining how the tensor produced by $\mathbf{T} \times \mathbf{U}$ is of $\operatorname{rank} k+(n-k+1)-2=n-1$.

There exists an isomorphism between tensors and multilinear maps (Bourbaki, 1989; Lee, 1997), such that any curried multilinear map

$$
f: V_{1} \rightarrow \ldots \rightarrow V_{j} \rightarrow V_{k}
$$

can be represented as a tensor $\mathbf{T}^{f} \in V_{k} \otimes V_{j} \otimes \ldots \otimes V_{1}$ (note the reversed order of the vector spaces), with tensor contraction acting as function application. This isomorphism guarantees that there exists such a tensor $\mathbf{T}^{f}$ for every $f$, such that the following equality holds for any $\mathbf{v}_{1} \in V_{1}, \ldots, \mathbf{v}_{j} \in V_{j}$ :

$$
f \mathbf{v}_{1} \ldots \mathbf{v}_{j}=\mathbf{v}_{k}=\mathbf{T}^{f} \times \mathbf{v}_{1} \times \ldots \times \mathbf{v}_{j}
$$

## 4 Tensor-based predicate calculi

In this section, I discuss how the isomorphism between multilinear maps and tensors described above can be used to model predicates, relations, and logical atoms of a predicate calculus. The four aspects of a predicate calculus we must replicate here using tensors are as follows: truth values, the logical domain and its elements (logical atoms), predicates, and relations. I will discuss logical connectives in the next section.

Both truth values and domain objects are the basic elements of a predicate calculus, and therefore it makes sense to model them as vectors rather than higher rank tensors, which I will reserve for relations. We first must consider the vector space used
to model the boolean truth values of $\mathbb{B}$. Coecke et al. (2010) suggest, as boolean vector space, the space $B$ with the basis $\{T, \perp\}$, where $T=\left[\begin{array}{ll}1 & 0\end{array}\right]^{\top}$ is interpreted as 'true', and $\perp=\left[\begin{array}{ll}0 & 1\end{array}\right]^{\top}$ as 'false'.

I assign to the domain $\mathcal{D}$, the set of objects in our logic, a vector space $D$ on $\mathbb{R}^{|\mathcal{D}|}$ with basis vectors $\left\{\mathbf{d}_{i}\right\}_{i}$ which are in bijective correspondence with elements of $\mathcal{D}$. An element of $\mathcal{D}$ is therefore represented as a one-hot vector in $D$, the single nonnull value of which is the weight for the basis vector mapped to that element of $\mathcal{D}$. Similarly, a subset of $\mathcal{D}$ is a vector of $D$ where those elements of $\mathcal{D}$ in the subset have 1 as their corresponding basis weights in the vector, and those not in the subset have 0 . Therefore there is a one-to-one correspondence between the vectors in $D$ and the elements of the power set $\mathcal{P}(\mathcal{D})$, provided the basis weights of the vectors are restricted to one of 0 or 1 .

Each unary predicate $P$ in the logic is represented in the logical model as a set $M_{P} \subseteq \mathcal{D}$ containing the elements of the domain for which the predicate is true. Predicates can be viewed as a unary function $f_{P}: \mathcal{D} \rightarrow \mathbb{B}$ where

$$
f_{P}(x)= \begin{cases}\top & \text { if } x \in M_{P} \\ \perp & \text { otherwise }\end{cases}
$$

These predicate functions can be modelled as rank 2 tensors in $B \otimes D$, i.e. matrices. Such a matrix $\mathbf{M}^{P}$ is expressed in sum notation as follows:

$$
\mathbf{M}^{P}=\left(\sum_{i} c_{1 i}^{M^{P}} \mathrm{~T} \otimes \mathbf{d}_{i}\right)+\left(\sum_{i} c_{2 i}^{M^{P}} \perp \otimes \mathbf{d}_{i}\right)
$$

The basis weights are defined in terms of the set $M_{P}$ as follows: $c_{1 i}^{M^{P}}=1$ if the logical atom $x_{i}$ associated with basis weight $\mathbf{d}_{i}$ is in $M_{P}$, and 0 otherwise; conversely, $c_{2 i}^{M^{P}}=1$ if the logical atom $x_{i}$ associated with basis weight $\mathbf{d}_{i}$ is not in $M_{P}$, and 0 otherwise.

To give a simple example, let's consider a domain with three individuals, represented as the following one-hot vectors in $D:$ john $=\left[\begin{array}{lll}1 & 0 & 0\end{array}\right]^{\top}$, chris $=\left[\begin{array}{lll}0 & 1 & 0\end{array}\right]^{\top}$, and tom $=\left[\begin{array}{lll}0 & 0 & 1\end{array}\right]^{\top}$. Let's imagine that Chris and John are mathematicians, but Tom is not. The predicate $P$ for 'is a mathematician' therefore is represented model-theoretically as the set $M_{P}=\{$ chris, john $\}$. Translating this into a matrix gives the following tensor for $P$ :

$$
\mathbf{M}^{P}=\left[\begin{array}{lll}
1 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]
$$

To compute the truth value of 'John is a mathematician', we perform predicate-argument application as tensor contraction (matrix-vector multiplication, in this case):

$$
\mathbf{M}^{P} \times \mathbf{j o h n}=\left[\begin{array}{lll}
1 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]\left[\begin{array}{l}
0 \\
1 \\
0
\end{array}\right]=\left[\begin{array}{l}
1 \\
0
\end{array}\right]=\top
$$

Likewise for 'Tom is a mathematician':

$$
\mathbf{M}^{P} \times \mathbf{t o m}=\left[\begin{array}{lll}
1 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]\left[\begin{array}{l}
0 \\
0 \\
1
\end{array}\right]=\left[\begin{array}{l}
0 \\
1
\end{array}\right]=\perp
$$

Model theory for predicate calculus represents any $n$-ary relation $R$, such as a verb, as the set $M_{R}$ of $n$-tuples of elements from $\mathcal{D}$ for which $R$ holds. Therefore such relations can be viewed as functions $f_{R}: \mathcal{D}^{n} \rightarrow \mathbb{B}$ where:

$$
f_{R}\left(x_{1}, \ldots, x_{n}\right)= \begin{cases}\top & \text { if }\left(x_{1}, \ldots, x_{n}\right) \in M_{R} \\ \perp & \text { otherwise }\end{cases}
$$

We can represent the boolean function for such a relation $R$ as a tensor $\mathbf{T}^{R}$ in $B \otimes \underbrace{D \otimes \ldots \otimes D}_{n}$ :

$$
\begin{aligned}
\mathbf{T}^{R}= & \left(\sum_{\alpha_{1} \ldots \alpha_{n}} c_{1 \alpha_{1} \ldots \alpha_{n}}^{T^{R}} \top \otimes \mathbf{d}_{\alpha_{1}} \otimes \ldots \otimes \mathbf{d}_{\alpha_{n}}\right) \\
& +\left(\sum_{\alpha_{1} \ldots \alpha_{n}} c_{2 \alpha_{1} \ldots \alpha_{n}}^{T^{R}} \perp \otimes \mathbf{d}_{\alpha_{1}} \otimes \ldots \otimes \mathbf{d}_{\alpha_{n}}\right)
\end{aligned}
$$

As was the case for predicates, the weights for relational tensors are defined in terms of the set modelling the relation: $c_{1 \alpha_{1} \ldots \alpha_{n}}^{T^{R}}$ is 1 if the tuple $(x, \ldots, z)$ associated with the basis vectors $\mathbf{d}_{\alpha_{n}} \ldots \mathbf{d}_{\alpha_{1}}$ (again, note the reverse order) is in $M_{R}$ and 0 otherwise; and $c_{2 \alpha_{1} \ldots \alpha_{n}}^{T^{R}}$ is 1 if the tuple $(x, \ldots, z)$ associated with the basis vectors $\mathbf{d}_{\alpha_{n}} \ldots \mathbf{d}_{\alpha_{1}}$ is not in $M_{R}$ and 0 otherwise.

To give an example involving relations, let our domain be the individuals John ( $j$ ) and Mary ( $m$ ). Mary loves John and herself, but John only loves himself. The logical model for this scenario is as follows:

$$
\mathcal{D}=\{j, m\} \quad M_{\text {loves }}=\{(j, j),(m, m),(m, j)\}
$$

Distributionally speaking, the elements of the domain will be mapped to the following one-hot vectors in some two-dimensional space $D$ as follows:
$\mathbf{j}=\left[\begin{array}{ll}1 & 0\end{array}\right]^{\top}$ and $\mathbf{m}=\left[\begin{array}{ll}0 & 1\end{array}\right]^{\top}$. The tensor for 'loves' can be written as follows, ignoring basis elements with null-valued basis weights, and using the distributivity of the tensor product over addition:

$$
\begin{aligned}
\mathbf{T}^{\text {loves }} & =\top \otimes\left(\left(\mathbf{d}_{1} \otimes \mathbf{d}_{1}\right)+\left(\mathbf{d}_{2} \otimes \mathbf{d}_{2}\right)+\left(\mathbf{d}_{1} \otimes \mathbf{d}_{2}\right)\right) \\
& +\left(\perp \otimes \mathbf{d}_{2} \otimes \mathbf{d}_{1}\right)
\end{aligned}
$$

Computing "Mary loves John" would correspond to the following calculation:

$$
\begin{aligned}
& \left(\mathbf{T}^{\text {loves }} \times \mathbf{m}\right) \times \mathbf{j}= \\
& \quad\left(\left(\mathrm{T} \otimes \mathbf{d}_{2}\right)+\left(\mathrm{T} \otimes \mathbf{d}_{1}\right)\right) \times \mathbf{j}=\mathrm{T}
\end{aligned}
$$

whereas "John loves Mary" would correspond to the following calculation:

$$
\begin{aligned}
& \left(\mathbf{T}^{\text {loves }} \times \mathbf{j}\right) \times \mathbf{m}= \\
& \quad\left(\left(\mathrm{T} \otimes \mathbf{d}_{1}\right)+\left(\perp \otimes \mathbf{d}_{2}\right)\right) \times \mathbf{m}=\perp
\end{aligned}
$$

## 5 Logical connectives with tensors

In this section, I discuss how the boolean connectives of a propositional calculus can be modelled using tensors. Combined with the predicate and relation representations discussed above, these form a complete quantifier-free predicate calculus based on tensors and tensor contraction.

Negation has already been shown to be modelled in the boolean space described earlier by Coecke et al. (2010) as the swap matrix:

$$
\mathbf{T}\urcorner=\left[\begin{array}{ll}
0 & 1 \\
1 & 0
\end{array}\right]
$$

This can easily be verified:

$$
\begin{aligned}
& \mathbf{T}\urcorner \times \top=\left[\begin{array}{ll}
0 & 1 \\
1 & 0
\end{array}\right]\left[\begin{array}{l}
1 \\
0
\end{array}\right]=\left[\begin{array}{l}
0 \\
1
\end{array}\right]=\perp \\
& \mathbf{T} \times \perp=\left[\begin{array}{ll}
0 & 1 \\
1 & 0
\end{array}\right]\left[\begin{array}{l}
0 \\
1
\end{array}\right]=\left[\begin{array}{l}
1 \\
0
\end{array}\right]=\top
\end{aligned}
$$

All other logical operators are binary, and hence modelled as rank 3 tensors. To make talking about rank 3 tensors used to model binary operations easier, I will use the following block matrix notation for $2 \times 2 \times 2$ rank 3 tensors $\mathbf{T}$ :

$$
\mathbf{T}=\left[\begin{array}{ll|ll}
a_{1} & b_{1} & a_{2} & b_{2} \\
c_{1} & d_{1} & c_{2} & d_{2}
\end{array}\right]
$$

which allows us to express tensor contractions as follows:

$$
\begin{aligned}
\mathbf{T} \times \mathbf{v} & =\left[\begin{array}{ll|ll}
a_{1} & b_{1} & a_{2} & b_{2} \\
c_{1} & d_{1} & c_{2} & d_{2}
\end{array}\right]\left[\begin{array}{c}
\alpha \\
\beta
\end{array}\right] \\
& =\left[\begin{array}{ll}
\alpha \cdot a_{1}+\beta \cdot a_{2} & \alpha \cdot b_{1}+\beta \cdot b_{2} \\
\alpha \cdot c_{1}+\beta \cdot c_{2} & \alpha \cdot d_{1}+\beta \cdot d_{2}
\end{array}\right]
\end{aligned}
$$

or more concretely:

$$
\left.\begin{array}{l}
\mathbf{T} \times \mathrm{T}=\left[\begin{array}{ll|ll}
a_{1} & b_{1} & a_{2} & b_{2} \\
c_{1} & d_{1} & c_{2} & d_{2}
\end{array}\right]\left[\begin{array}{l}
1 \\
0
\end{array}\right]=\left[\begin{array}{ll}
a_{1} & b_{1} \\
c_{1} & d_{1}
\end{array}\right] \\
\mathbf{T} \times \perp=\left[\begin{array}{lll}
a_{1} & b_{1} & a_{2} \\
b_{2} \\
c_{1} & d_{1} & c_{2}
\end{array} d_{2}\right.
\end{array}\right]\left[\begin{array}{l}
0 \\
1
\end{array}\right]=\left[\begin{array}{ll}
a_{2} & b_{2} \\
c_{2} & d_{2}
\end{array}\right] .
$$

Using this notation, we can define tensors for the following operations:

$$
\begin{aligned}
& (\vee) \mapsto \mathbf{T}^{\vee}=\left[\begin{array}{ll|ll}
1 & 1 & 1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right] \\
& (\wedge) \mapsto \mathbf{T}^{\wedge}=\left[\begin{array}{ll|ll}
1 & 0 & 0 & 0 \\
0 & 1 & 1 & 1
\end{array}\right] \\
& (\rightarrow) \mapsto \mathbf{T}^{\rightarrow}=\left[\begin{array}{ll|ll}
1 & 0 & 1 & 1 \\
0 & 1 & 0 & 0
\end{array}\right]
\end{aligned}
$$

I leave the trivial proof by exhaustion that these fit the bill to the reader.

It is worth noting here that these tensors preserve normalised probabilities of truth. Let us consider a model such at that described in Coecke et al. (2010) which, in lieu of boolean truth values, represents truth value vectors of the form $[\alpha \beta]^{\top}$ where $\alpha+\beta=1$. Applying the above logical operations to such vectors produces vectors with the same normalisation property. This is due to the fact that the columns of the component matrices are all normalised (i.e. each column sums to 1). To give an example with conjunction, let $\mathbf{v}=\left[\alpha_{1} \beta_{1}\right]^{\top}$ and $\mathbf{w}=\left[\alpha_{2} \beta_{2}\right]^{\top}$ with $\alpha_{1}+\beta_{1}=\alpha_{2}+\beta_{2}=1$. The conjunction of these vectors is calculated as follows:

$$
\begin{aligned}
& \left(\mathbf{T}^{\wedge} \times \mathbf{v}\right) \times \mathbf{w} \\
& \quad=\left[\begin{array}{ll|ll}
1 & 0 & 0 & 0 \\
0 & 1 & 1 & 1
\end{array}\right]\left[\begin{array}{c}
\alpha_{1} \\
\beta_{1}
\end{array}\right]\left[\begin{array}{c}
\alpha_{2} \\
\beta_{2}
\end{array}\right] \\
& \quad=\left[\begin{array}{cc}
\alpha_{1} & 0 \\
\beta_{1} & \alpha_{1}+\beta_{1}
\end{array}\right]\left[\begin{array}{c}
\alpha_{2} \\
\beta_{2}
\end{array}\right] \\
& \quad=\left[\begin{array}{c}
\alpha_{1} \alpha_{2} \\
\beta_{1} \alpha_{2}+\left(\alpha_{1}+\beta_{1}\right) \beta_{2}
\end{array}\right]
\end{aligned}
$$

To check that the probabilities are normalised we calculate:

$$
\begin{aligned}
& \alpha_{1} \alpha_{2}+\beta_{1} \alpha_{2}+\left(\alpha_{1}+\beta_{1}\right) \beta_{2} \\
& =\left(\alpha_{1}+\beta_{1}\right) \alpha_{2}+\left(\alpha_{1}+\beta_{1}\right) \beta_{2} \\
& =\left(\alpha_{1}+\beta_{1}\right)\left(\alpha_{2}+\beta_{2}\right)=1
\end{aligned}
$$

We can observe that the resulting probability distribution for truth is still normalised. The same property can be verified for the other connectives, which I leave as an exercise for the reader.

## 6 Quantifiers and non-linearity

The predicate calculus described up until this point has repeatedly been qualified as 'quantifier-free', for the simple reason that quantification cannot be modelled if each application of a predicate or relation immediately yields a truth value. In performing such reductions, we throw away the information required for quantification, namely the information which indicates which elements of a domain the predicate holds true or false for. In this section, I present a variant of the predicate calculus developed earlier in this paper which allows us to model simple quantification (i.e. excluding embedded quantifiers) alongside a tensor-based approach to predicates. However, I will prove that this approach to quantifier modelling relies on non-linear functions, rendering them non-suitable for compositional distributional models relying solely on multilinear maps for composition (or alternatively, rendering such models unsuitable for the modelling of quantifiers by this method).

We saw, in Section 4, that vectors in the semantic space $D$ standing for the logical domain could model logical atoms as well as sets of atoms. With this in mind, instead of modelling a predicate $P$ as a truth-function, let us now view it as standing for some function $f_{P}: \mathcal{P}(\mathcal{D}) \rightarrow \mathcal{P}(\mathcal{D})$, defined as:

$$
f_{P}(X)=X \cap M_{P}
$$

where $X$ is a set of domain objects, and $M_{P}$ is the set modelling the predicate. The tensor form of such a function will be some $\mathbf{T}^{f_{P}}$ in $D \otimes D$. Let this square matrix be a diagonal matrix such that basis weights $c_{i i}^{T_{f_{p}}}=1$ if the atom $x$ corresponding to $\mathbf{d}_{i}$ is in $M_{P}$ and 0 otherwise. Through tensor contraction, this
tensor maps subsets of $\mathcal{D}$ (elements of $D$ ) to subsets of $\mathcal{D}$ containing only those objects of the original subset for which $P$ holds (i.e. yielding another vector in $D$ ).

To give an example: let us consider a domain with two dogs ( $a$ and $b$ ) and a cat ( $c$ ). One of the dogs ( $b$ ) is brown, as is the cat. Let $S$ be the set of dogs, and $P$ the predicate "brown". I represent these statements in the model as follows:

$$
\mathcal{D}=\{a, b, c\} \quad S=\{a, b\} \quad M_{P}=\{b, c\}
$$

The set of dogs is represented as a vector $\mathbf{S}=$ $\left[\begin{array}{lll}1 & 1 & 0\end{array}\right]^{\top}$ and the predicate 'brown' as a tensor in $D \otimes D:$

$$
\mathbf{T}^{P}=\left[\begin{array}{lll}
0 & 0 & 0 \\
0 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]
$$

The set of brown dogs is obtained by computing $f_{B}(S)$, which distributionally corresponds to applying the tensor $\mathbf{T}^{P}$ to the vector representation of $S$ via tensor contraction, as follows:

$$
\mathbf{T}^{P} \times \mathbf{S}=\left[\begin{array}{lll}
0 & 0 & 0 \\
0 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]\left[\begin{array}{l}
1 \\
1 \\
0
\end{array}\right]=\left[\begin{array}{l}
0 \\
1 \\
0
\end{array}\right]=\mathbf{b}
$$

The result of this computation shows that the set of brown dogs is the singleton set containing the only brown dog, $b$. As for how logical connectives fit into this picture, in both approaches discussed below, conjunction and disjunction are modelled using set-theoretic intersection and union, which are simply the component-wise min and max functions over vectors, respectively.

Using this new way of modelling predicates as tensors, I turn to the problem of modelling quantification. We begin by putting all predicates in vector form by replacing each instance of the bound variable with a vector 1 filled with ones, which extracts the diagonal from the predicate matrix.

An intuitive way of modelling universal quantification is as follows: expressions of the form "All $X \mathrm{~s}$ are $Y \mathrm{~s}$ " are true if and only if $M_{X}=M_{X} \cap M_{Y}$, where $M_{X}$ and $M_{Y}$ are the set of $X \mathrm{~s}$ and the set of $Y \mathrm{~s}$, respectively. Using this, we can define the map forall for distributional universal quantification modelling expressions of the form "All $X \mathrm{~s}$ are $Y \mathrm{~s}$ " as follows:

$$
\operatorname{forall}(\mathbf{X}, \mathbf{Y})= \begin{cases}\top & \text { if } \mathbf{X}=\min (\mathbf{X}, \mathbf{Y}) \\ \perp & \text { otherwise }\end{cases}
$$

To give a short example, the sentence 'All Greeks are human' is verified by computing $\mathbf{X}=\left(\mathbf{M}^{\text {greek }} \times \mathbf{1}\right)$, $\mathbf{Y}=\left(\mathbf{M}^{\text {human }} \times \mathbf{1}\right)$, and verifying the equality $\mathbf{X}=$ $\min (\mathbf{X}, \mathbf{Y})$.

Existential statements of the form "There exists X" can be modelled using the function exists, which tests whether or not $M_{X}$ is empty, and is defined as follows:

$$
\operatorname{exists}(\mathbf{X})= \begin{cases}\top & \text { if }|\mathbf{X}|>0 \\ \perp & \text { otherwise }\end{cases}
$$

To give a short example, the sentence 'there exists a brown dog' is verified by computing $\mathbf{X}=\left(\mathbf{M}^{\text {brown }} \times\right.$ 1) $\cap\left(\mathbf{M}^{\operatorname{dog}} \times \mathbf{1}\right)$ and verifying whether or not $\mathbf{X}$ is of strictly positive length.

An important point to note here is that neither of these quantification functions are multi-linear maps, since a multilinear map must be linear in all arguments. A counter example for forall is to consider the case where $M_{X}$ and $M_{Y}$ are empty, and multiply their vector representations by non-zero scalar weights $\alpha$ and $\beta$.

$$
\begin{aligned}
& \alpha \mathbf{X}=\mathbf{X} \\
& \beta \mathbf{Y}=\mathbf{Y} \\
& \operatorname{forall}(\alpha \mathbf{X}, \beta \mathbf{Y})=\operatorname{forall}(\mathbf{X}, \mathbf{Y})=\top \\
& \operatorname{forall}(\alpha \mathbf{X}, \beta \mathbf{Y}) \neq \alpha \beta \top
\end{aligned}
$$

I observe that the equations above demonstrate that forall is not a multilinear map.

The proof that exists is not a multilinear map is equally trivial. Assume $M_{X}$ is an empty set and $\alpha$ is a non-zero scalar weight:

$$
\begin{aligned}
& \alpha \mathbf{X}=\mathbf{X} \\
& \operatorname{exists}(\alpha \mathbf{X})=\operatorname{exists}(\mathbf{X})=\perp \\
& \operatorname{exists}(\alpha \mathbf{X}) \neq \alpha \perp
\end{aligned}
$$

It follows that exists is not a multi-linear function.

## 7 Conclusions and future work

In this paper, I set out to demonstrate that it was possible to replicate most aspects of predicate logic using tensor-based models. I showed that tensors can be constructed from logical models to represent predicates and relations, with vectors encoding elements or sets of elements from the logical domain.

I discussed how tensor contraction allows for evaluation of logical expressions encoded as tensors, and that logical connectives can be defined as tensors to form a full quantifier-free predicate calculus. I exposed some of the limitations of this approach when dealing with variables under the scope of quantifiers, and proposed a variant for the tensor representation of predicates which allows us to deal with quantification. Further work on tensor-based modelling of quantifiers should ideally seek to reconcile this work with that of Barwise and Cooper (1981). In this section, I discuss how both of these approaches to predicate modelling can be put into relation, and suggest further work that might be done on this topic, and on the topic of integrating this work into compositional distributional models of semantics.

The first approach to predicate modelling treats predicates as truth functions represented as tensors, while the second treats them as functions from subsets of the domain to subsets of the domain. Yet both representations of predicates contain the same information. Let $\mathbf{M}^{P}$ and $\mathbf{M}^{P}$ be the tensor representations of a predicate $P$ under the first and second approach, respectively. The relation between these representations lies in the equality $\operatorname{diag}\left(\mathbf{p} \mathbf{M}^{P}\right)=$ $\mathbf{M}^{\prime \mathbf{P}}$, where $\mathbf{p}$ is the covector [10] (and hence $\mathbf{p} \mathbf{M}^{P}$ yields the first row of $\mathbf{M}^{P}$ ). The second row of $\mathbf{M}^{P}$ being defined in terms of the first, one can also recover $\mathbf{M}^{P}$ from the diagonal of $\mathbf{M}^{P}$.

Furthermore, both approaches deal with separate aspects of predicate logic, namely applying predicates to logical atoms, and applying them to bound variables. With this in mind, it is possible to see how both approaches can be used sequentially by noting that tensor contraction allows for partial application of relations to logical atoms. For example, applying a binary relation to its first argument under the first tensor-based model yields a predicate. Translating this predicate into the second model's form using the equality defined above then permits us to use it in quantified expressions. Using this, we can evaluate expressions of the form "There exists someone who John loves". Future work in this area should therefore focus on developing a version of this tensor calculus which permits seamless transition between both tensor formulations of logical predicates.

Finally, this paper aims to provide a starting point for the integration of logical aspects into composi-
tional distributional semantic models. The work presented here serves to illustrate how tensors can simulate logical elements and operations, but does not address (or seek to address) the fact that the vectors and matrices in most compositional distributional semantic models do not cleanly represent elements of a logical domain. However, such distributional representations can arguably be seen as representing the properties objects of a logical domain hold in a corpus: for example the similar distributions of 'car' and 'automobile' could serve to indicate that these concepts are co-extensive. This suggests two directions research based on this paper could take. One could use the hypothesis that similar vectors indicate co-extensive concepts to infer a (probabilistic) logical domain and set of predicates, and use the methods described above without modification; alternatively one could use the form of the logical operations and predicate tensors described in this paper as a basis for a higher-dimensional predicate calculus, and investigate how such higher-dimensional 'logical' operations and elements could be defined or learned. Either way, the problem of reconciling the fuzzy 'messiness' of distributional models with the sharp 'cleanliness' of logic is a difficult problem, but I hope to have demonstrated in this paper that a small step has been made in the right direction.

## Acknowledgments

Thanks to Ondřej Rypáček, Nal Kalchbrenner and Karl Moritz Hermann for their helpful comments during discussions surrounding this paper. This work is supported by EPSRC Project EP/I03808X/1.

## References

M. Baroni and R. Zamparelli. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, pages 1183-1193. Association for Computational Linguistics, 2010.
J. Barwise and R. Cooper Generalized quantifiers and natural language. Linguistics and philosophy, pages 159-219. Springer, 1981.
W. Blacoe and M. Lapata. A comparison of vector-based representations for semantic composition. Proceed-
ings of the 2012 Conference on Empirical Methods in Natural Language Processing, 2012.
N. Bourbaki. Commutative Algebra: Chapters 1-7. Springer-Verlag (Berlin and New York), 1989.
S. Clark and S. Pulman. Combining symbolic and distributional models of meaning. In AAAI Spring Symposium on Quantum Interaction, 2006.
B. Coecke, M. Sadrzadeh, and S. Clark. Mathematical Foundations for a Compositional Distributional Model of Meaning. Linguistic Analysis, volume 36, pages 345-384. March 2010.
J. R. Curran. From distributional to semantic similarity. PhD thesis, 2004.
D. R. Dowty, R. E. Wall, and S. Peters. Introduction to Montague Semantics. Dordrecht, 1981.
J. R. Firth. A synopsis of linguistic theory 1930-1955. Studies in linguistic analysis, 1957.
M. Fitting. First-order logic and automated theorem proving. Springer Verlag, 1996.
E. Grefenstette, G. Dinu, Y. Zhang, M. Sadrzadeh, and M. Baroni. Multi-step regression learning for compositional distributional semantics. In Proceedings of the Tenth International Conference on Computational Semantics. Association for Computational Linguistics, 2013.
E. Grefenstette and M. Sadrzadeh. Experimental support for a categorical compositional distributional model of meaning. In Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, 2011.
G. Grefenstette. Explorations in automatic thesaurus discovery. 1994.
D. Kartsaklis, and M. Sadrzadeh and S. Pulman. A Unified Sentence Space for Categorical DistributionalCompositional Semantics: Theory and Experiments. In Proceedings of 24th International Conference on Computational Linguistics (COLING 2012): Posters, 2012.
T. K. Landauer and S. T. Dumais. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. Psychological review, 1997.
J. Lee. Riemannian manifolds: An introduction to curvature, volume 176. Springer Verlag, 1997.
D. W. Loveland. Automated theorem proving: A logical basis. Elsevier North-Holland, 1978.
J. Mitchell and M. Lapata. Vector-based models of semantic composition. In Proceedings of ACL, volume 8, 2008.
J. Mitchell and M. Lapata. Language models based on semantic composition. In Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1, pages 430-439. Association for Computational Linguistics, 2009.
J. J. Mitchell. Composition in distributional models of semantics. PhD thesis, 2011.
R. Montague. English as a Formal Language. Formal Semantics: The Essential Readings, 1974.
H. Schütze. Automatic word sense discrimination. Computational linguistics, 24(1):97-123, 1998.
P. Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist systems. Artificial intelligence, 46(1-2):159-216, 1990.
R. Socher, B. Huval, C.D. Manning, and A.Y Ng. Semantic compositionality through recursive matrixvector spaces. Proceedings of the 2012 Conference on Empirical Methods in Natural Language Processing, pages 1201-1211, 2012.
P. D. Turney. Domain and function: A dual-space model of semantic relations and compositions. Journal of Artificial Intelligence Research, 44:533-585, 2012.
F. M. Zanzotto, I. Korkontzelos, F. Fallucchi, and S. Manandhar. Estimating linear models for compositional distributional semantics. In Proceedings of the 23 rd International Conference on Computational Linguistics, pages 1263-1271. Association for Computational Linguistics, 2010.

