

# RELPRON: A Relative Clause Evaluation Dataset for Compositional Distributional Semantics

Laura Rimell\*  
University of Cambridge Computer  
Laboratory

Jean Maillard\*  
University of Cambridge Computer  
Laboratory

Tamara Polajnar\*  
University of Cambridge Computer  
Laboratory

Stephen Clark\*  
University of Cambridge Computer  
Laboratory

*This article introduces RELPRON, a large dataset of subject and object relative clauses, for the evaluation of methods in compositional distributional semantics. RELPRON targets an intermediate level of grammatical complexity between content-word pairs and full sentences. The task involves matching terms, e.g. 'wisdom', with representative properties, e.g. 'quality that experience teaches'. A unique feature of RELPRON is that it is built from attested properties, but without the need for them to appear in relative clause format in the source corpus. The article also presents some initial experiments on RELPRON, using a variety of composition methods, from simple baselines, through to arithmetic operators on vectors, and finally more complex methods in which argument-taking words are represented as tensors. The latter methods are based on the Categorical framework, which is described in detail. The results show that vector addition is difficult to beat — in line with the existing literature — but that an implementation of the Categorical framework based on the Practical Lexical Function model is able to match the performance of vector addition. The article finishes with an in-depth analysis of RELPRON, showing how results vary across subject and object relative clauses, across the different head nouns, and how the methods perform on the subtasks necessary for capturing relative clause semantics, as well as a qualitative analysis highlighting some of the more common errors. Our hope is that the competitive results presented here, in which the best systems are on average ranking one out of every two properties correctly for a given term, will inspire new approaches to the RELPRON ranking task and other linguistically interesting constructions.*

## 1. Introduction

The field of compositional distributional semantics (Mitchell and Lapata 2008; Clark, Coecke, and Sadrzadeh 2008; Baroni, Bernardi, and Zamparelli 2014) integrates distributional semantic representations of words (Schütze 1998; Turney and Pantel 2010; Clark 2015) with formal methods for composing word representations into larger phrases and sentences (Montague 1974; Dowty, Wall, and Peters 1981). In recent years a number of

---

\* University of Cambridge Computer Laboratory, William Gates Building, 15 JJ Thomson Avenue, Cambridge, UK. E-mail: first.last@cl.cam.ac.uk.

composition methods have been proposed, including simple arithmetic operations on distributional word vectors (Mitchell and Lapata 2008, 2010), multi-linear operations involving higher-order representations of argument-taking words such as verbs and adjectives (Baroni and Zamparelli 2010; Coecke, Sadrzadeh, and Clark 2010), and composition of distributed word vectors learned with neural networks (Socher, Manning, and Ng 2010; Mikolov, Yih, and Zweig 2013). To compare such approaches it is important to have high-quality datasets for evaluating composed phrase representations at different levels of granularity and complexity.

Existing evaluation datasets fall largely into two categories. Some datasets focus on two-to-three-word phrases consisting of content words only (i.e. no closed-class function words), including subject-verb, verb-object, subject-verb-object, and adjective-noun combinations (Mitchell and Lapata 2008, 2010; Baroni and Zamparelli 2010; Vecchi, Baroni, and Zamparelli 2011; Grefenstette and Sadrzadeh 2011; Boleda et al. 2012, 2013; Lazaridou, Vecchi, and Baroni 2013). Such datasets have been essential for evaluating the first generation of compositional distributional models, but the relative grammatical simplicity of the phrases — in particular, the absence of function words — leaves open a wide range of compositional phenomena in natural language against which models must be evaluated.<sup>1</sup>

Other datasets, including those used in recent SemEval and \*SEM Shared Tasks (Agirre et al. 2012, 2013; Marelli et al. 2014; Agirre 2015), focus on pairs of full sentences, some as long as 20 words or more. The evaluation is based on a numeric similarity measure for each sentence pair. These datasets represent a more realistic task for language technology applications, but reducing sentence similarity to a single value makes it difficult to identify how a model performs on subparts of a sentence, or on specific grammatical constructions, masking the areas where models need improvement.<sup>2</sup>

In the domain of syntactic parsing, a call has been made for “grammatical construction-focused” parser evaluation (Rimell, Clark, and Steedman 2009; Bender et al. 2011), focusing on individual, often challenging syntactic structures, in order to tease out parser performance in these areas from overall accuracy scores. We make an analogous call here, for a wider range of compositional phenomena to be investigated in compositional distributional semantics in the near future.

This article begins to answer that call by presenting RELPRON, a dataset of noun phrases consisting of a noun modified by a relative clause. For example, *building that hosts premieres* is a noun phrase containing a subject relative clause (a relative clause with an extracted subject), and in our dataset describes a *theater*; while *person that helicopter saves* contains an object relative clause, and in our dataset describes a *survivor*. The RELPRON dataset is primarily concerned with the analysis of a particular type of closed-class function word, namely relative pronouns (*that*, in our examples); function words have traditionally been of greater concern than content words for formal semantics, and we address how this focus can be extended to distributional semantics.

Relative clauses, while still fairly short and therefore more manageable than full sentences, are nevertheless more grammatically complex than the short phrases in previous datasets, due to the relative pronoun and the long-distance relationship between the verb and the extracted argument, known as the head noun of the relative clause (*building*, *person*). The aim of RELPRON is to expand the variety of phrase types on which

---

1 Exceptions are the datasets of Bernardi et al. (2013) and Pham et al. (2013), which include determiners such as *two*, *most*, and *no* in the phrases to be composed; see Section 2.1.

2 The pilot subtask on interpretable Semantic Textual Similarity (Agirre 2015) begins to address this problem.

compositional distributional semantic methods can be evaluated, thus helping to build these methods up step-by-step to the full task of compositional sentence representations.

RELPRON is a large (1,087 relative clauses), corpus-based, naturalistic dataset, including both subject and object relative clauses that modify a variety of concrete and abstract nouns. The relative clauses are matched with terms and represent distinctive properties of those terms, such as *wisdom: quality that experience teaches*, and *bowler: player that dominates batsman*. A unique feature of RELPRON is that it is built from attested properties of terms, but without the properties needing to appear in relative clause form in the source corpus.

This article also presents some initial experiments on RELPRON, using a variety of composition methods. We find that a simple arithmetic vector operation provides a challenging baseline, in line with existing literature evaluating such methods, but we are able to match this baseline using a more sophisticated method similar to the Practical Lexical Function (PLF) model of Paperno, Pham, and Baroni (2014). We hope that the compositional methods presented here will inspire new approaches.

In the remainder of the article, Section 2 provides some motivation for the RELPRON dataset and the canonical ranking task that we imagine RELPRON being used for, as well as a description of existing datasets designed to test compositional distributional models. Section 3 describes the dataset itself, and provides a detailed description of how it was built. Section 4 surveys some of the existing compositional methods that have been applied in distributional semantics, including a short historical section describing work from cognitive science in the 1980s and 90s. Section 4.2 provides a fairly detailed description of the Categorical framework, which provides the basis for the more complex composition methods that we have tested, as well as a description of existing implementations of the framework. Section 5 provides further details of the composition methods, including lower-level details of how the vectors and tensors were built. Section 6 presents our results on the development and test portions of RELPRON, comparing different composition methods and different methods of building the vectors and tensors (including count-based vectors vs. neural embeddings). Finally, Section 7 provides an in-depth analysis of RELPRON, showing how results vary across subject and object relative clauses, across the different head nouns, and how the methods perform on the subtasks necessary for capturing relative pronoun semantics, as well as providing a qualitative analysis highlighting some of the more common errors.

## 2. Motivation for RELPRON

Relative clauses have been suggested as an interesting test case for compositional distributional semantic methods by Sadrzadeh, Clark, and Coecke (2013) and Baroni, Bernardi, and Zamparelli (2014). One of the main inspirations for RELPRON is the pilot experiment of Sadrzadeh, Clark, and Coecke (2013), in which terms such as *carnivore* are matched with descriptions such as *animal who eats meat*. However, the dataset is small and all examples are manually constructed. RELPRON aims to provide a larger, more realistic, and more challenging testing ground for relative clause composition methods.

RELPRON falls into the general category of *cross-level semantic similarity* tasks (Jurgens, Pilehvar, and Navigli 2014), which involve comparing phrases of different lengths; in this case, a single word with a short phrase. The task of matching a phrase with a single word can be thought of as a special case of paraphrase recognition. However, restricting one of the phrases to a word means that the experiment is better controlled with regard to composition methods; since we know that we have available high-quality,

state-of-the-art distributional representations of single words, the only unknown is the composed phrase representation, whereas a comparison of two composed phrases can be more difficult to interpret.

The RELPRON task is similar to the definition classification task of Kartsaklis, Sadrzadeh, and Pulman (2012), which targeted verb phrase composition. In that task, methods were tested on their ability to match a term (in this case a verb) with its definition (usually of the form verb-object), e.g. *embark*: *enter boat or vessel*. However, the definitions of Kartsaklis, Sadrzadeh, and Pulman (2012) were mined from a set of dictionaries, whereas the RELPRON relative clauses are not limited to dictionary definitions.

Because there is an intuitive definition-like “flavour” to some relative clauses, we considered using dictionary definitions as a source of relative clauses for the dataset. However, in practice we found that short, natural definitions in relative clause format are rare in dictionaries. For example, definitions for the word *telescope* from three different dictionaries, two of them learner dictionaries, are:

<i>telescope</i>	a monocular optical instrument possessing magnification for observing distant objects (Wiktionary)
<i>telescope</i>	an important tool for astronomy that gathers and focuses light (Simple Wikipedia)
<i>telescope</i>	a cylinder-shaped device for making objects that are far away look nearer and larger, using lenses and curved mirrors (Cambridge Advanced Learner’s Dictionary)


Only the Simple Wikipedia definition uses a relative clause to define the term, and would still require significant editing to shorten the phrase. Instead, we collect relative clauses that express representative properties, not necessarily definitions, of the term, such as:

<i>telescope</i>	device that astronomers use
<i>telescope</i>	device that detects planets
<i>telescope</i>	device that observatory has

Collecting properties rather than definitions allows for several phrases to be associated with each term, and substantially broadens the number and type of corpora which can be used as sources for the relative clauses.

Figure 1 introduces some terminology for describing the terms and properties in the dataset.<sup>3</sup> A *term* is a single word which is associated with a set of descriptive phrases. Semantically, each descriptive phrase is a *property* of the term. Syntactically, each descriptive phrase is a *noun phrase containing a relative clause*. For the sake of brevity, we often refer to the whole noun phrase containing the relative clause as a *relative clause*, when we believe this will not lead to confusion. We say *subject relative clause* or *object relative clause* (or *subject property* or *object property*) to indicate which argument of the verb has been extracted. The *head noun* of the relative clause is the extracted noun, and is a hypernym of the term; for example, *device* is a hypernym of *telescope*. The subject or object which remains in situ in the clause we call the *argument* of the verb, and the general term for the subject or object role is *grammatical function*. The relative clauses in RELPRON are restrictive relative clauses, because they narrow down the meaning of the

<sup>3</sup> For readability, we sometimes present verbs and nouns in their inflected form in relative clauses in this article, e.g. *detects planets* rather than *detect planet*. In the dataset, all words are lemmatised.

**Figure 1**

Terminology used to describe the terms and properties in RELPRON: subject relative clause above, object relative clause below.

head noun. In contrast, a non-restrictive relative clause provides incidental situational information, e.g. *a device, which was in the room yesterday*.

Our goal was to construct a dataset containing at least 1,000 properties in the form of relative clauses. To make the ranking task more challenging, we chose groups of terms which share a head noun: for example, *telescope*, *watch*, *button*, and *pipe* are all subclasses of *device*, so their relative clauses will all begin *device that...* Since the relative clauses are not definitions, each term can be associated with multiple properties; we selected between four and ten properties per term.

The construction of RELPRON also takes a novel approach to producing evaluation datasets that target a particular grammatical construction. Rather than mining examples of the construction directly from a corpus, which may be difficult due to data sparsity, the properties in RELPRON did not have to occur in relative clause format in our source corpus. Instead, the properties were obtained from subject-verb-object (SVO) triples such as *astronomer uses telescope*, and joined with head nouns such as *device*. The head nouns are also attested in the corpus as heads of relative clauses in general, though not necessarily with these specific properties. The relative clauses in RELPRON are therefore *naturalistic*, *plausible*, and *corpus-based*, without having to be attested in the exact relative clause form in which they ultimately appear in the dataset.

The canonical task on the RELPRON dataset is conceived as a ranking task. Given a term, a system must rank all properties by their similarity to the term. The ideal system will rank all properties corresponding to that term above all other properties. This is analogous to an Information Retrieval (IR) task (Manning, Raghavan, and Schütze 2008) where the term is the query and the properties are the documents, with the identifying properties for a given term as the relevant documents. Evaluation can make use of standard IR measures; in this article we use Mean Average Precision (MAP).

## 2.1 Existing Evaluation Datasets

Most existing evaluation datasets for compositional distributional semantics have been based on small, fixed syntactic contexts, typically adjective-noun or subject-verb-object. Mitchell and Lapata (2008) introduce an evaluation which has been adopted by a number of other researchers. The task is to predict human similarity judgements on subject-verb phrases, where the similarity rating depends on word sense disambiguation in context. For example, *horse run* is more similar to *horse gallop* than to *horse dissolve*, while *colors run* is more similar to *colors dissolve* than to *colors gallop*. Compositional models are evaluated on how well their similarity judgments correlate with the human judgments. Mitchell and Lapata (2010) introduce a phrase similarity dataset consisting

of adjective-noun, verb-object, and noun-noun combinations, where the task again is to produce similarity ratings that correlate well with human judgements. For example, *large number* and *great majority* have high similarity in the gold-standard data, while *further evidence* and *low cost* have low similarity. This dataset formed the basis of the shared task at the GEMS 2011 workshop (Padó and Peirsman 2011). Grefenstette and Sadrzadeh (2011) and Kartsaklis and Sadrzadeh (2014) introduced analogous tasks for subject-verb-object triples. In the disambiguation task, *people try door* is more similar to *people test door* than to *people judge door*. In the similarity task, *medication achieve result* and *drug produce effect* have high similarity ratings, while *author write book* and *delegate buy land* have low ratings. Sentence pairs with mid-similarity ratings tend to have high relatedness but are not mutually substitutable, e.g. *team win match* and *people play game*.

Other evaluations have made use of the contrast between compositional and non-compositional phrases, based on the intuition that a successful compositional model should produce phrase vectors similar to the observed context vectors for compositional phrases, but diverge from the observed context vectors for phrases known to be non-compositional. Reddy, McCarthy, and Manandhar (2011) introduce an evaluation based on compound nouns like *climate change* (compositional) and *gravy train* (non-compositional). Boleda et al. (2012, 2013) evaluate models on a dataset involving intersective adjectives in phrases like *white towel* (compositional) and non-intersective adjectives in phrases like *black hole* and *false floor* (non-compositional). The data from the shared task of the ACL 2011 Distributional Semantics and Compositionality workshop (Biemann and Giesbrecht 2011) is also relevant in this context.

Also relevant for compositional distributional semantics are larger datasets for full sentence similarity and entailment, including those from recent SemEval and \*SEM Shared Tasks (Agirre et al. 2012, 2013; Marelli et al. 2014; Agirre 2015). The SemEval Semantic Textual Similarity (STS) tasks make use of a number of paraphrase corpora which include text from news articles, headlines, tweets, image captions, video descriptions, student work, online discussion forums, and statistical machine translation. These corpora include sentences of widely varying length and quality. The SICK dataset (Marelli et al. 2014) is based on image and video description corpora, with named entities removed and other normalization steps applied, in order to remove phenomena outside the scope of compositional distributional semantic models. Additional sentences are added based on linguistic transformations which vary the meaning of the original sentence in a controlled fashion. However, this dataset is not designed for targeted evaluation of models on specific compositional linguistic phenomena, especially since many of the sentence pairs exhibit significant lexical overlap, e.g. *An airplane is in the air*, *An airplane is flying in the air*.

Some recent datasets have extended the range of linguistic phenomena which may be evaluated. Cheung and Penn (2012) test the ability of compositional models to detect semantic similarity in the presence of lexical and syntactic variation, for example among the sentences *Pfizer buys Rinat*, *Pfizer paid several hundred million dollars for Rinat*, and *Pfizer's acquisition of Rinat*. Bernardi et al. (2013) introduce a dataset comparing nouns to noun phrases, including determiners as function words; for example, *duel* is compared to *two opponents* (similar), *various opponents* (dissimilar), and *two engineers* (dissimilar). Pham et al. (2013) introduce a dataset involving full sentences from a limited grammar, where determiners and word order vary; for example, *A man plays a guitar* is compared to *A guitar plays a man* and *The man plays a guitar*.

### 3. The RELPRON Dataset

One challenge in targeting specific grammatical constructions for evaluation is that the more linguistically interesting constructions can be relatively rare in text corpora. The development of RELPRON takes a novel approach, which uses a syntactic transformation to assemble properties in relative clause form without the properties needing to appear in this syntactic form in the source corpus. The properties were obtained from SVO triples such as *astronomer uses telescope*, and joined with head nouns such as *device*, a hypernym of *telescope*, to create the relative clause *device that astronomer uses*. The source corpus used for constructing the dataset was the union of a 2010 Wikipedia download and the British National Corpus (BNC) (Burnard 2007). Here we describe the steps in the construction of RELPRON.

#### 3.1 Selection of Candidate Head Nouns

To make the dataset maximally challenging, we designed it with multiple terms (hyponyms) per head noun (hypernym). This ensures that the composition method cannot rely solely on the similarity between the term and the head noun — say, *traveler* and *person* — and fail to consider the verb and argument in the relative clause.

We began by selecting a set of candidate head nouns, each of which would need to have a number of good candidate terms. To improve the naturalness of the examples, we chose candidate head nouns from among nouns frequently used as heads of relative clauses in the source corpus. We used a simple regular expression search for nouns occurring in the pattern “a[n] NOUN that|who|which”, considering as candidates the 200 most frequent results. We manually excluded nouns where the *that*-clause could be an argument of the noun, for example *statement*, *decision*, and *belief*.

One author (Rimell) used WordNet (Miller 1995) as a guide to manually narrow down the candidates, with the following two requirements. First, the head nouns should have a large number of WordNet hyponyms, to maximise the chances of finding appropriate terms. The hyponyms were required to be direct children or grandchildren, since we wanted to avoid extremely abstract nouns like *entity*, which has many distant hyponyms but few direct ones. Second, the hyponyms should be mostly unigrams, because we wanted to avoid multiword terms such as *remote control*. For simplicity, we considered only the first WordNet sense of each candidate head noun.<sup>4</sup> Occasionally we utilised the WordNet level one down from the frequent head noun in the corpus; for example, the corpus search yielded *substance*, but we used its direct hyponym *material* as a head noun, because it had more direct hyponyms of its own.

#### 3.2 Selection of Candidate Terms

Following the selection of candidate head nouns, we chose candidate terms. For each candidate head noun, one author (Rimell) manually selected candidate terms from among the full list of WordNet hyponyms. Only one word was chosen from each WordNet synset, on the grounds that members of the same synset would not be sufficiently distinguishable from one another by their properties. Generally this was the

---

<sup>4</sup> For e.g. *person*, the WordNet command for viewing the hyponym hierarchy was: `wn person -n1 -treen`.

**Table 1**

Sample SVO triples extracted from the source corpus for the candidate term *traveler*.

Object	Subject
it allow traveler	traveler make decision
website allow traveler	traveler make journey
inn serve traveler	traveler take road
business serve traveler	traveler take advantage
she become traveler	traveler have option
student become traveler	traveler have luggage

first unigram member of the synset. For example, the WordNet hyponyms for *building* include the synset:

dormitory, dorm, residence hall, hall, student residence

From this synset *dormitory* was kept as a candidate term. On the other hand, the WordNet hyponyms for *quality* include the synset:

economy, thriftiness

Here *economy* was discarded on the basis that its primary sense is not relevant, and *thriftiness* was kept as a candidate term. Entire synsets were discarded when they seemed to be a poor example of the head noun, to be too ambiguous, or consisted only of multi-word or proper noun entries. A multi-word term was very occasionally edited to make a unigram.

Next, potential terms were filtered by frequency. Based on initial exploration of the data, we passed along for manual annotation only those candidate terms that occurred at least 5,000 times in the source corpus, to maximise the chances that there would be enough SVO triples for annotation (see Section 3.3). Finally, candidate terms that appeared in WordNet as hyponyms of more than one candidate head noun were removed. For example, *surgery* could be an *activity* or a *room*; *school* could be a *building* or an *organization*. We believed that the inclusion of such terms would complicate the dataset and the evaluation unnecessarily.

### 3.3 Extraction of SVO Triples

For each candidate term, candidate properties were obtained from the source corpus by extraction of all triples in which the term occurred in either subject or object position. We used a version of the corpus lemmatised with Morpha (Minnen, Carroll, and Pearce 2001) and parsed with the C&C Tools (Clark and Curran 2007) to obtain the subject and object relations. A minimum frequency cutoff of 6 occurrences per verb, per grammatical function (subject/object), per candidate term was applied. Properties containing proper noun arguments were removed. Table 1 shows some sample triples extracted from the corpus prior to manual annotation.

### 3.4 Conversion to Relative Clause Form

The candidate properties were converted to relative clause form prior to annotation, by joining the head noun with the relevant part of the SVO triple, using the template *term: headnoun that V O* for a subject property, or *term: headnoun that S V* for an object property. Using the examples from Table 1, *inn serve traveler* thus became *traveler: person that inn serve*; and *traveler make journey* became *traveler: person that make journey*. For simplicity, *that* was used as the relative pronoun throughout the dataset, although some head nouns are more likely to appear with *who* or *which*.

### 3.5 Manual Annotation of Properties

The manual annotation step was the stage at which candidate head nouns, terms, and properties were filtered for the final dataset. Only terms with at least four good properties, and head nouns with at least four terms, were retained after this step. Annotation was continued until the dataset reached the desired size.

Most of the SVO triples extracted from the corpus did not make good identifying properties for their terms. In many cases the SVO triple represented a property which, while it could be true of the term, did not distinguish it from other terms, especially from other hyponyms of the same head noun. Thus *player that wins tournament* is not a good property for *golfer*: many golfers do win tournaments, but so do other sports players. Similarly, *organization that money helps* is an accurate description of a *charity*, but does not distinguish it from other types of organization. In other cases, the extracted triple contained a light verb and/or a pronoun in argument position, such as *golfer: player that he become*, or *charity: organization that these include*. These candidates were too vague to be meaningful properties. Some candidate properties were unsuitable because they were not generic. Among the candidates for *killer* were *person that win lottery* and *person that actor play*, because some specific killer in the corpus experienced these events. However, they are not descriptive of killers in general. Lexical ambiguity and parser errors contributed other unsuitable candidates.

Two of the authors (Rimell and Clark) selected good properties from among the candidates for each term, looking for what we called *identifying properties*.<sup>5</sup> Such a property should distinguish a term from other hyponyms of its head noun, both those in RELPRON, and ideally other examples of the head noun as well (barring near-synonyms such as *traveler* and *voyager*). As a further guideline for annotation, we defined an identifying property as one that distinguishes a term because it *mainly* or *canonically* applies to the term in question, compared with other examples of the head noun. For example, *person that tavern serves* was not judged to be a good property for *traveler*. It does not distinguish a *traveler* from other examples of people, because serving *travelers* is not the main or canonical purpose of a tavern (since the majority of people who eat and drink in a tavern are not traveling). On the other hand, *person that hotel accommodates* is a good property, because, while hotels might accommodate other people (e.g. for meetings or in a restaurant), their main and canonical purpose is to accommodate *travelers*. A more subtle example is the property *device that bomb has*, a candidate property for *timer*. This property does not uniquely distinguish *timer* from among other devices, since bombs contain multiple (sub-)devices, but we felt that a timer was such a canonical part of a bomb that this could be considered a good property for *timer*.

<sup>5</sup> Full annotation guidelines will be made publicly available along with the corpus.

**Table 2**Full list of selected properties for the terms *navy* and *expert* from the development set.

OBJ	navy: organization that sailor join	OBJ	expert: person that novice become
OBJ	navy: organization that fleet destroy	OBJ	expert: person that panel include
OBJ	navy: organization that vessel serve	OBJ	expert: person that lawyer consult
OBJ	navy: organization that battleship fight	OBJ	expert: person that report cite
SBJ	navy: organization that use submarine	SBJ	expert: person that have knowledge
SBJ	navy: organization that maintain blockade	SBJ	expert: person that give tutorial
SBJ	navy: organization that defeat fleet	SBJ	expert: person that describe model
SBJ	navy: organization that protect waters	SBJ	expert: person that write treatise
SBJ	navy: organization that establish blockade	SBJ	expert: person that develop curriculum
SBJ	navy: organization that blockade port	SBJ	expert: person that author monograph

For a given head noun (hypernym), the annotator began with a random selection of ten candidate terms (hyponyms) from the candidates selected as in Section 3.2. For each candidate term, the annotator chose up to ten good properties. Once ten had been found, annotation stopped on that term, even if more good properties might have been available, because the aim was a broad and balanced dataset rather than recall of all good properties. Not all terms had ten good properties. We required a minimum of four, otherwise the candidate term was discarded and another candidate randomly selected from the list of candidates. We also required a minimum of four terms per head noun, otherwise the entire head noun was discarded and another chosen from the list of candidates.

The two annotators originally worked together on a small set of data. After the annotation guidelines were refined, the head nouns were divided between the annotators. As a final check, all properties were reviewed and discussed; any property on which the annotators could not agree was discarded. We aimed for a balance of subject and object relative clauses for each term, but this was not always possible. Some head nouns were very unbalanced in this respect: *person* terms tended to have better subject properties due to their agentivity, e.g. *philosopher: person that interprets world*; while *quality* terms tended to have better object properties due to their lack of agentivity, e.g. *popularity: quality that artist achieves*. We also tried to make the dataset more challenging by choosing properties with lexical confounders when we noticed them. For example, *team: organization that loses match* and *division: organization that team joins* share the lemma *team*, which means a ranking method cannot rely on lexical overlap between the term and the property as an indicator of similarity. Table 2 shows the full list of selected properties for two terms in the development set, *navy* and *expert*.

Finding terms with enough good properties was a challenging task. Overall, we had to examine candidate properties for 284 candidate terms to yield the 138 terms in the final dataset. Most head nouns in the final dataset have ten terms, but two (*player* and *woman*) have only five terms, while one (*scientist*) has eight terms. The final list of head nouns, in alphabetical order, is: *activity, building, device, document, mammal, material, organization, person, phenomenon, player, quality, room, scientist, vehicle, woman*.<sup>6</sup>

<sup>6</sup> Readers may wonder why *woman* and not *man* was included in the dataset. This was a simple matter of statistics: *man* did not have enough WordNet hyponyms with sufficient corpus frequency, according to the criteria in Section 3.2.

**Table 3**

Total number of terms and properties by head noun in RELPRON.

Test Set			Development Set		
Head noun	Terms	Props	Head noun	Terms	Props
phenomenon	10	80	quality	10	81
activity	10	83	organization	10	99
material	10	82	device	10	76
room	10	80	building	10	69
vehicle	10	79	document	10	69
mammal	10	70	person	10	86
woman	5	37	player	5	38
scientist	8	58			
TOTAL	73	569	TOTAL	65	518

**Table 4**

Total number of subject and object properties by head noun in RELPRON.

Test Set			Development Set		
Head noun	Sbj	Obj	Head noun	Sbj	Obj
phenomenon	38	42	quality	16	65
activity	46	37	organization	54	45
material	30	52	device	40	36
room	38	42	building	37	32
vehicle	38	41	document	21	48
mammal	41	29	person	59	27
woman	20	17	player	29	9
scientist	58	0			
TOTAL	309	260	TOTAL	256	262

### 3.6 Final Dataset

After annotation was complete, the data was divided into test and development sets. The division was made by head noun (hypernym), with all terms (and therefore properties) for a given head noun appearing in either the test or the development set. Assuming that terms with the same head noun are good mutual confounders, this split was chosen to maximise the difficulty of the task. We arranged the split so that both sets contained a range of concrete and abstract head nouns. We also ensured that both sets contained the maximum possible number of lexical confounders.

The final dataset consists of 1,087 properties, comprising 15 head nouns and 138 terms. This is divided into a test set of 569 properties, comprising 8 head nouns and 73 terms; and a development set of 518 properties, comprising 7 head nouns and 65 terms. The total number of terms and properties by head noun is given in Table 3. Overall, the dataset includes 565 subject and 522 object properties, with the distribution across test and development sets as shown in Table 4.

#### 4. Composition Methods – Background

This section surveys some of the composition methods that have been proposed in the distributional semantics literature, focusing on the methods applied in this article. We also include a short section giving more of a historical perspective, contrasting current work in computational linguistics with that from cognitive science in the 1980s and 90s. Although in our experiments we do use neural embeddings (Mikolov, Yih, and Zweig 2013) as input vectors to the composition, we leave the application of neural networks to the modelling of the relative pronoun itself as important future work. Hence in this section we leave out work using recurrent or recursive neural networks specifically designed for composition, e.g. (Socher, Manning, and Ng 2010; Socher et al. 2013; Hermann, Grefenstette, and Blunsom 2013).

In this section, we are agnostic about how the vector representations are built. Many of the composition techniques we describe can be equally applied to “classical” distributional vectors built using count methods (Turney and Pantel 2010; Clark 2015) and vectors built using neural embedding techniques. It may be that some composition techniques work better with one type of vector or the other — for example there is evidence that elementwise multiplication performs better with count methods — although research comparing the two vector-building methods is still in its infancy (Baroni, Dinu, and Kruszewski 2014; Milajevs et al. 2014; Levy and Goldberg 2014; Levy et al. 2015). Section 5 gives details of how our vector representations are built in practice.

##### 4.1 Simple Operators on Vectors

The simplest method of combining two distributional representations, for example vectors for *fast* and *car*, is to perform some elementwise combination, such as vector addition or elementwise multiplication. As a method of integrating formal and distributional semantics, these operators appear to be clearly inadequate *a priori*, not least because they are commutative, and so do not respect word order. Despite this fact, there is a large literature investigating how such operators can be used for phrasal composition, starting with the work of Mitchell and Lapata (2008) and Mitchell and Lapata (2010) (M&L subsequently).

The additive model from M&L has the general form:

$$\mathbf{p} = A\mathbf{u} + B\mathbf{v} \quad (1)$$

where  $\mathbf{u}$  and  $\mathbf{v}$  are word (column) vectors in  $\mathbb{R}^n$  (e.g.  $\vec{\text{fast}}$  and  $\vec{\text{car}}$ ),  $\mathbf{p} \in \mathbb{R}^n$  is the vector for the phrase resulting from composition of  $\mathbf{u}$  and  $\mathbf{v}$  ( $\vec{\text{fast car}}$ ), and  $A \in \mathbb{R}^{n \times n}$  and  $B \in \mathbb{R}^{n \times n}$  are matrices which determine the contribution of  $\mathbf{u}$  and  $\mathbf{v}$  to  $\mathbf{p}$ . However, M&L make the simplifying assumption that only the  $i$ th components of  $\mathbf{u}$  and  $\mathbf{v}$  contribute to the  $i$ th component of  $\mathbf{p}$ , which yields the form:

$$\mathbf{p} = \alpha\mathbf{u} + \beta\mathbf{v} \quad (2)$$

The parameters  $\alpha, \beta \in \mathbb{R}$  allow the contributions of  $\mathbf{u}$  and  $\mathbf{v}$  to be weighted differently, providing a minimal level of syntax-awareness to the model. For example, if  $\mathbf{u}$  is the vector representation for an adjective, and  $\mathbf{v}$  the vector representation for a noun,  $\mathbf{v}$  may

be given a higher weight since it provides the linguistic head for the resulting noun phrase; whereas if  $\mathbf{u}$  is an intransitive verb and  $\mathbf{v}$  a noun,  $\mathbf{u}$  may be weighted more highly.

The multiplicative model presented by M&L has the general form:

$$\mathbf{p} = C(\mathbf{u} \otimes \mathbf{v}) \quad (3)$$

where  $\mathbf{u} \otimes \mathbf{v}$  is the tensor, or outer, product of  $\mathbf{u}$  and  $\mathbf{v}$ , and  $C$  is a tensor of order three, which projects the tensor product of  $\mathbf{u}$  and  $\mathbf{v}$  onto the space of  $\mathbf{p}$  (which is assumed to be the same space as that containing  $\mathbf{u}$  and  $\mathbf{v}$ ). Under similar simplifying assumptions to the additive model, the multiplicative model has the form:

$$\mathbf{p} = \mathbf{u} \odot \mathbf{v} \quad (4)$$

where  $\odot$  represents elementwise multiplication, that is,  $\mathbf{p}_i = \mathbf{u}_i \cdot \mathbf{v}_i$ , where  $\mathbf{p}_i$  is the  $i$ th coefficient of  $\mathbf{p}$ . In this simpler model,  $C$  simply picks out the diagonal in the  $\mathbf{u} \otimes \mathbf{v}$  tensor.

One obvious question to consider is how well such simple operators scale when applied to phrases or sentences longer than a few words. Polajnar, Rimell, and Clark (2014) and Polajnar and Clark (2014) suggest that elementwise multiplication performs badly in this respect, with the quality of the composed representation degrading quickly as the number of composition operations increases. Vector addition is more stable, but Polajnar, Rimell, and Clark (2014) suggest that the quality of the composed representation degrades after around ten binary composition operations, even with addition.

M&L also introduce a method that aims to interpret one of the constituent vectors as an operator, despite its lying in the same semantic space as the argument. A dilation operation decomposes  $\mathbf{v}$  (the argument) into two components, one parallel to  $\mathbf{u}$  and the other perpendicular, and dilates  $\mathbf{v}$  in the direction of  $\mathbf{u}$ :

$$\mathbf{p} = (\mathbf{u} \cdot \mathbf{u})\mathbf{v} + (\lambda - 1)(\mathbf{u} \cdot \mathbf{v})\mathbf{u} \quad (5)$$

M&L also consider the tensor product, which is an instance of the general multiplicative model described above (when  $C$  is the identity operator), and described below in Section 4.4, but find that it performs poorly on their phrasal composition tasks and datasets, using classical count vectors. One of the potential problems with the tensor product is that the resulting phrase vector lives in a different vector space to the argument vectors, so that  $\vec{fast} \otimes \vec{car}$ , for example, lives in a different space to  $\vec{car}$ . Hence M&L also discuss circular convolution, a technique for projecting a vector in a tensor product space onto the original vector space components (Plate 1991).

## 4.2 The Categorical Framework

Since the majority of the methods used in our experiments are variants of the Categorical framework, this section provides a fairly detailed description of that framework.

In an attempt to unite a formal, Montague-style semantics (Montague 1974; Dowty, Wall, and Peters 1981) with a distributional semantics, Coecke, Sadrzadeh, and Clark

$$\overrightarrow{red} \quad \overrightarrow{car} \quad \overrightarrow{red\ car}$$

$$\begin{pmatrix} R_{11} & R_{12} & R_{13} & R_{14} & R_{15} \\ R_{21} & R_{22} & R_{23} & R_{24} & R_{25} \\ R_{31} & R_{32} & R_{33} & R_{34} & R_{35} \\ R_{41} & R_{42} & R_{43} & R_{44} & R_{45} \\ R_{51} & R_{52} & R_{53} & R_{54} & R_{55} \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \\ c_5 \end{pmatrix} = \begin{pmatrix} rc_1 \\ rc_2 \\ rc_3 \\ rc_4 \\ rc_5 \end{pmatrix}$$

**Figure 2**

Matrix multiplication for adjective-noun combinations.

(2010) and Baroni, Bernardi, and Zamparelli (2014) both make the observation that the semantics of argument-taking words such as verbs and adjectives can be represented using multi-linear algebraic objects.<sup>7</sup> This is the basis for the Categorical framework (Coecke, Sadrzadeh, and Clark 2010), which uses a generalisation of matrices, i.e. higher-order tensors, for words which take more than one argument. The Categorical framework has vectors and their multi-linear analogues as native elements, and provides a natural operation for composition, namely tensor contraction (see Section 4.2.2). It therefore stands in contrast to frameworks for compositional distributional semantics that effectively take the logic of formal semantics as a starting point, and use distributional semantics to enrich the logical representations (Lewis and Steedman 2013; Garrette, Erk, and Mooney 2011; Herbelot and Copestake 2015).

It is important to note that the Categorical framework makes no commitment to how the multi-linear algebraic objects are realised, only a commitment to their shape and how they are combined. In particular, the vector spaces corresponding to atomic categories, such as noun and sentence, do not have to be *distributional* in the classical sense; they could be *distributed* in the connectionist sense (Smolensky 1990), where the basis vectors themselves are not readily interpretable (the neural embeddings that we use in this article fall into this category).

**4.2.1 Composition as Matrix Multiplication.** A useful starting point in describing the framework is adjectival modification. In fact, the proposal in Baroni and Zamparelli (2010) (B&Z hereafter) to model the meanings of adjectives as matrices can be seen as an instance of the Categorical framework (and also an instance of the general framework described in Baroni, Bernardi, and Zamparelli (2014)). The insight in B&Z is that, in theoretical linguistics, adjectives are typically thought of as having a functional role, mapping noun denotations to noun denotations. B&Z make the transition to vector spaces by arguing that, in linear algebra, functions are represented as matrices (the *linear maps*). This insight provides an obvious answer to the question of what the composition operator should be in this framework, namely matrix multiplication. Figure 2 shows how the context vector for *car* is multiplied by the matrix for *red* to give the vector

<sup>7</sup> Another way to motivate this idea is that some words naturally have more of an “operator semantics”, whereas others have more of a “content semantics”. Socher et al. (2012, 2013) realize this distinction in the context of a recursive neural network, using matrices for the operator semantics and vectors for the content semantics. Every word has both associated types, but the network may learn to put more weight on one type or the other for a given word.

for *red car*. In this simple example, the noun space containing  $\vec{car} = (c_1, \dots, c_5)^T$  and  $\vec{red\ car} = (rc_1, \dots, rc_5)^T$  has 5 basis vectors, which means that a  $5 \times 5$  matrix,  $\overline{red}$ , is required for the adjective.

The second contribution of B&Z is to propose a method for learning the  $\overline{red}$  matrix from supervised training data. What should the gold-standard representation for *red car* be? B&Z argue that, given large enough corpora, it should ideally be the context vector for the compound *red car*; i.e. that the vector for *red car* should be built in exactly the same way as the vector for *car*. These context vectors for phrases are generally known as *holistic vectors*. The reason for the adjective matrix is that it allows the generalisation of adjective-noun combinations beyond those seen in the training data. The details of the B&Z training process will not be covered in this section, but briefly the process is to find all examples of adjective-noun pairs in a large training corpus, and then use standard linear regression techniques to obtain matrices for each adjective in the training data. These learnt matrices can then be used to generalise beyond the seen adjective-noun pairs. So if *tasty artichoke*, for example, has not been seen in the training data (or perhaps seen infrequently), the prediction for the context vector *tasty artichoke* can be obtained, as long as there are enough examples of *tasty X* in the data to obtain the  $\overline{tasty}$  matrix (via linear regression), and enough occurrences of *artichoke* to obtain the  $\overline{artichoke}$  vector.

Testing is performed by using held-out context vectors for some of the adjective-noun pairs in the data, and seeing how close — according to the cosine measure — the predicted context vector is to the actual context vector for each adjective-noun pair in the test set. The learning method using linear regression was compared against various methods of combining the context vectors of the adjective and noun, such as vector addition and elementwise multiplication, and was found to perform significantly better at this prediction task.

**4.2.2 An Extension Using Multi-Linear Algebra: Tensor-Based CCG semantics.** The grammatical formalism assumed in this subsection will be a variant of Categorical Grammar, which was also the grammar formalism used by Montague (1974). The original papers setting out the tensor-based compositional framework (Clark, Coecke, and Sadrzadeh 2008; Coecke, Sadrzadeh, and Clark 2010) used pregroup categorial grammars (Lambek 2008), largely because they share an abstract mathematical structure with vector spaces. However, other forms of categorial grammar can be used equally as well in practice, and here we use Combinatory Categorical Grammar (CCG, Steedman (2000)), which has been shown to fit seamlessly with the tensor-based semantic framework (Grefenstette 2013; Maillard, Clark, and Grefenstette 2014). With a computational linguistics audience in mind, we also present the framework entirely in terms of multi-linear algebra, rather than the category theory of the original papers.

A matrix is a second-order tensor; for example, the  $\overline{red}$  matrix above lives in the  $N \otimes N$  space, meaning that two indices — each corresponding to a basis vector in the noun space  $N$  — are needed to specify an entry in the matrix.<sup>8</sup> Noting that  $N \otimes N$  is structurally similar to the CCG syntactic type for an adjective ( $N/N$ ) — both specify functions from nouns to nouns — a recipe for translating a syntactic type into a semantic type suggests itself: replace all slash operators in the syntactic type with tensor

<sup>8</sup> The tensor product  $V \otimes W$  of two vector spaces  $V \in \mathbb{R}^n$  and  $W \in \mathbb{R}^m$  is an  $nm$ -dimensional space spanned by elements of the form  $\vec{v} \otimes \vec{w}$ , i.e. pairs of basis vectors of  $V$  and  $W$ . An element  $a \in V \otimes W$  can be written as  $\sum a_{ij} \vec{v}_i \otimes \vec{w}_j$  where  $a_{ij}$  is a scalar.

product operators. With this translation, the combinators employed by CCG to combine syntactic categories carry over seamlessly to the meaning spaces, maintaining what is often described as CCG’s “transparent interface” between syntax and semantics. The seamless integration with CCG arises from the (somewhat trivial) observation that tensors are linear maps — a particular kind of function — and hence can be manipulated using CCG’s combinatory rules.

Here are some example syntactic types, and the corresponding tensor spaces in which words with those types are semantically represented (using the notation *syntactic type* : *semantic type*). We first assume that all atomic types have meanings living in distinct vector spaces:<sup>9</sup>

- noun phrases,  $NP : N$
- sentences,  $S : S$

Replacing each slash in a complex syntactic type with a tensor product operator results in the following meaning spaces, following the CCG result-leftmost convention for both the syntactic and semantic types:

- Intransitive verb,  $S \backslash NP : S \otimes N$
- Transitive verb,  $(S \backslash NP) / NP : S \otimes N \otimes N$
- Ditransitive verb,  $((S \backslash NP) / NP) / NP : S \otimes N \otimes N \otimes N$
- Adverbial modifier,  $(S \backslash NP) \backslash (S \backslash NP) : S \otimes N \otimes S \otimes N$
- Preposition modifying  $NP$ ,  $(NP \backslash NP) / NP : N \otimes N \otimes N$

Hence the meaning of an intransitive verb, for example, is a particular matrix, or second-order tensor, in the tensor product space  $S \otimes N$ . The meaning of a transitive verb is a “cuboid”, or third-order tensor, in the tensor product space  $S \otimes N \otimes N$ . In the same way that the syntactic type of an intransitive verb can be thought of as a function — taking an  $NP$  and returning an  $S$  — the meaning of an intransitive verb is also a function (linear map) — taking a vector in  $N$  and returning a vector in  $S$ . Another way to think of this function is that each element of the matrix specifies, for a pair of basis vectors (one from  $N$  and one from  $S$ ), how a value on the given  $N$  basis vector contributes to the result for the given  $S$  basis vector.

In order to see how the tensors combine with their arguments, consider the following CCG derivation involving an intransitive verb, with the corresponding semantic type below its syntactic type:

$$\begin{array}{c} \text{Basil} \quad \text{sleeps} \\ \hline \begin{array}{cc} NP & S \backslash NP \\ N & S \otimes N \end{array} \\ \hline \begin{array}{c} S \\ S \end{array} < \end{array}$$

Let the meaning of *Basil* be  $\overrightarrow{\text{Basil}} \in N$  and *sleeps* be  $\overline{\text{sleeps}} \in S \otimes N$ ; then applying the  $\overline{\text{sleeps}}$  operator (matrix) to its argument  $\overrightarrow{\text{Basil}}$  is an example of the more general *tensor contraction* operation from multi-linear algebra. We refer readers to Maillard, Clark, and Grefenstette (2014) for details of how the additional combinators of CCG — e.g. forward

---

<sup>9</sup> We make no distinction in this section between nouns and noun phrases, as far as the corresponding semantic space is concerned.

and backward composition, backward-crossed composition, type-raising — naturally apply to the tensor-based semantic representations, and how tensor contraction applies in this more general setting.

*4.2.3 Existing Implementations of the Categorical Framework.* There are a few existing implementations of the Categorical framework, focusing mostly on adjectives and transitive verbs. The adjective implementation is that of B&Z as described in Section 4.2.1, where linear regression is used to learn each adjective as a mapping from noun contexts to adjective-noun contexts, with the observed context vectors for the noun and adjective-noun instances as training data. Since an adjective-noun combination has noun phrase meaning, the noun space is the obvious choice for the space in which the composed meanings should live. Maillard and Clark (2015) extend this idea by showing how the skip-gram model of Mikolov, Yih, and Zweig (2013) can be adapted to learn adjective matrices as part of a neural embedding objective.

For verb-argument composition, the question of sentence spaces arises.<sup>10</sup> A transitive verb such as *chase* lives in  $S \otimes N \otimes N$ , and can therefore be represented as

$$\sum_{ijk} C_{ijk} (\vec{s}_i \otimes \vec{n}_j \otimes \vec{n}_k) \quad (6)$$

for some choice of coefficients  $C_{ijk}$ . Grefenstette et al. (2011) choose  $S = N \otimes N$ , using the tensor product of the noun space with itself as the sentence space. Their concrete implementation learns the verb by counting observed  $n^{(s)}, n^{(o)}$  pairs where  $n^{(s)}$  is the subject and  $n^{(o)}$  the object of  $V$ . The final representation for  $V$  is then

$$V = \sum_{i=1}^{N_V} \vec{n}_i^{(s)} \otimes \vec{n}_i^{(o)} \quad (7)$$

where  $N_V$  is the number of instances of  $V$  in the corpus, and  $(n_i^{(s)}, n_i^{(o)})$  are the subject and object nouns for the  $i$ th instance of  $V$ . This method of learning the verb tensor is known as the *relational* method, since it captures relations between subject and object features. Composition with a particular subject and object reduces to taking the tensor product of the subject and object, and performing elementwise multiplication with the verb matrix, yielding a sentence meaning in  $N \otimes N$ . Sentence meanings can be compared to one another by taking the cosine of their angle in sentence space. One problem with this approach is that sentences with transitive verbs live in a different space than sentences with intransitive verbs, so sentences with different grammatical structures cannot be directly compared.

Kartsaklis, Sadrzadeh, and Pulman (2012) introduce another sentence space, by setting  $S = N$ . To achieve this, while training the verb in the same way as the relational method, they embed the  $N \otimes N$  tensor into  $N \otimes N \otimes N$  space, by copying the subject or object slices of the original tensor; for example, in the copy-object method the

<sup>10</sup> We consider a verb with all of its argument positions saturated, for example an SVO triple, to be a sentence, although real-world sentences are much longer and contain determiners and other function words.

representation for  $V$  is

$$V = \sum_{i=1}^{N_V} \vec{n}_i^{(s)} \otimes \vec{n}_i^{(o)} \otimes \vec{n}_i^{(o)} \quad (8)$$

Representing sentence meaning in the noun space confers the advantage that words, phrases, and sentences of any grammatical structure can be compared with one another, but it is an open question whether the features in the noun space, whether distributional or not, are appropriate or adequate for representing sentence meaning.

An alternative way of learning the transitive verb tensor parameters is presented in Grefenstette et al. (2013), using a process analogous to B&Z's process for learning adjectives. Two linear regression steps are performed. In the first step, a matrix is learned representing verb-object phrases, that is, a verb that has already been paired with its object. For example, the matrix for *eat meat* is learned as a mapping from corpus instances such as *dogs* and *dogs eat meat*. The full tensor for *eat* is then learned with *meat* as input and the *eat meat* matrix as output. Essentially, the subject mapping is learned in the first step, and the object mapping in the second.

Polajnar, Fagarasan, and Clark (2014), following Krishnamurthy and Mitchell (2013), investigate a non-distributional sentence space, the "plausibility space" described by Clark (2013, 2015). Here the sentence space is one- or two-dimensional, with sentence meaning either a real number between 0 and 1, or a probability distribution over the classes *plausible* and *implausible*. A verb tensor is learned using single-step linear regression, with the training data consisting of positive (plausible) and negative (implausible) SVO examples. Positive examples are attested in the corpus, and negative examples for a given verb have frequency-matched random nouns substituted in the argument positions. One tensor investigated has dimensions  $K \times K \times S$ , where  $K$  is the number of noun dimensions and  $S$  has two dimensions, each ranging from 0 to 1. The parameters of the tensor are learned so that when combined with a subject and object, a plausibility judgement is produced. Polajnar, Rimell, and Clark (2015) investigate a distributional sentence space based on the wider discourse context, in which the meaning of a sentence is represented as a distributional vector based on context words in the surrounding discourse, making the sentence space completely distinct from the noun space.<sup>11</sup>

Existing implementations of the Categorical framework have investigated learning up to third-order tensors (for transitive verbs). However, in practice, syntactic categories such as  $((N/N)/(N/N))/((N/N)/(N/N))$  are not uncommon in the wide-coverage CCG grammar of Hockenmaier and Steedman (2007); such a category would require an *eighth-order* tensor. The combination of many word-category pairs and higher-order tensors results in a huge number of parameters to be learned in any implementation. As a solution to this problem, various ways of reducing the number of parameters are being investigated, for example using tensor decomposition techniques (Kolda and Bader 2009; Fried, Polajnar, and Clark 2015), and removing some of the interactions encoded in the tensor by using only matrices to encode each predicate-argument combination (Paperno, Pham, and Baroni 2014; Polajnar, Fagarasan, and Clark 2014).

For this article, the problem of large numbers of parameters arises in the modelling of the relative pronoun, since its CCG type is  $(NP \backslash NP)/(S \backslash NP)$  for subject

<sup>11</sup> Kiros et al. (2015) investigate a similar sentence space within the context of recurrent neural networks.

relative clauses or  $(NP \setminus NP)/(S/NP)$  for object relative clauses, resulting in a fourth-order tensor,  $N \otimes N \otimes S \otimes N$ . Given the limited amount of training data available (see Section 5.5.1), we do not attempt to model the full tensor, but investigate two approximations, based on the methods of Paperno, Pham, and Baroni (2014). First, by modelling the transitive verb with two matrices — one for the subject interaction and the other for the object — we are able to model the verb as a pair of matrices with semantic type  $S \otimes N$ . This approximation is described in Section 5.2.2, and allows us to reduce the relative pronoun to a third-order tensor,  $N \otimes N \otimes S$  (Section 5.5.1). Second, we apply the same “decoupling” approximation to the relative pronoun, in order to represent the relative pronoun as a pair of matrices (Section 5.5.4), which can be learned independently of one another.

### 4.3 Relative Clause Composition

Two approaches to modelling relative clauses have been proposed in the compositional distributional semantics literature, both within the Categorical framework, although neither one has received a large-scale implementation. Both approaches are based on the fact that the relative clause (interpreted strictly, without the head noun) is a noun modifier, so the relative pronoun must map the meaning of the composed verb-argument phrase to a modifier of the head noun. In set theoretic terms, the relative pronoun signals the intersection between the set of individuals denoted by the head noun, and those denoted by the verb-argument phrase. For example, in *accuracy: quality that correction improves*, *accuracy* is both a quality and a thing that is improved by correction.

Clark, Coecke, and Sadrzadeh (2013) and Sadrzadeh, Clark, and Coecke (2013) propose an interpretation of relative clauses based on Frobenius algebra. The transitive verb is represented by a third-order tensor which combines with its argument by tensor contraction. The relative pronoun introduces a Frobenius operator which “discards” the third dimension of the verb tensor, allowing information to flow from the composed verb and argument to the head noun. In practice, this operator allows the verb to be represented as a matrix. The relative clause therefore takes the form of a vector which is then composed with the head noun ( $\vec{n}$ ) by elementwise multiplication:<sup>12</sup>

$$\vec{n} \odot (\bar{V} \vec{o}) \quad (\text{subject relative clause, with verb } V \text{ and object } o) \quad (9)$$

$$\vec{n} \odot (\bar{V}^T \vec{s}) \quad (\text{object relative clause, with verb } V \text{ and subject } s) \quad (10)$$

This method corresponds to an intuition that elementwise multiplication represents intersection, emphasizing those features which are shared by the head noun and the composed verb-argument phrase. As described in Section 2, this method has been tested on a toy dataset by Sadrzadeh, Clark, and Coecke (2013), using relational verb matrices; we test it on RELPRON in this article.

<sup>12</sup> In practice, the formulae work out equivalently to the copy-subject and copy-object methods used by Kartsaklis, Sadrzadeh, and Pulman (2012) for composed SVO triples, although they are arrived at in a different way. The formulation for the object relative clause also glosses over the fact that the CCG derivation requires type-raising and function composition, but it can easily be shown that the resulting semantic interpretation is equivalent based on the type-raising implementation in Maillard, Clark, and Grefenstette (2014).

Baroni, Bernardi, and Zamparelli (2014) propose a method for learning a relative pronoun as a fourth-order tensor which maps verb-argument phrases to noun modifiers. The proposed method would use multi-step linear regression to learn: (1) in the first step, verb-argument matrices, such as a *chase cats* matrix from training data such as  $\langle \text{dogs}, \text{dogs chase cats} \rangle$ ; (2) in the second step, noun modifier matrices, such as a *that chase cats* matrix from training data such as  $\langle \text{animal}, \text{animal that chases cats} \rangle$ ; and (3) in the third step, a tensor for the relative pronoun *that*, using the matrices from steps 1 and 2 with matched predicates. The final tensor would thus be learned from paired matrices such as *chases cats, that chases cats*. Under this approach, the intersective semantics of the relative pronoun must be captured within the tensor. Although the necessary training data might be relatively sparse, particularly for holistic relative clause vectors, Baroni, Bernardi, and Zamparelli (2014) hypothesize that they are sufficiently common to learn a general representation for the relative pronoun. We test a similar approach in this article, using holistic relative clause vectors to learn a tensor representing the relative pronoun (Section 5.5.1), though with single-step linear regression and a simplification of the verb representation that allows the tensor to be third- rather than fourth-order. We also learn a variant that represents the relative pronoun as a pair of matrices (Section 5.5.4) and requires learning fewer parameters.

#### 4.4 Historical Perspective

The general question of how to combine vector-based meaning representations in a compositional framework is a relatively new question for computational linguistics, but one which has been actively researched in cognitive science since the 1980s. The question arose because of the perceived failure of distributed models in general, and connectionist models in particular, to provide a suitable account of compositionality in language (Fodor and Pylyshyn 1988). The question is also relevant to the broad enterprise of Artificial Intelligence in that connectionist or distributed representations may seem to be in opposition to the more traditional symbolic systems often employed in AI.

Smolensky argues for the tensor product as the operation which binds a predicate to an argument in the distributed meaning space (Smolensky 1990; Smolensky and Legendre 2006). Smolensky represents a structure as possessing a set of *roles*,  $\mathbf{r}_i$ , which, for a particular instance of the structure, also has a set of *fillers*,  $\mathbf{f}_i$  (the roles are *bound* to fillers). The vector representation of a role-filler pair is obtained using the tensor product, and the representation for a complete structure,  $\mathbf{s}$ , of role-filler pairs is the sum of the tensor products for each pair:

$$\mathbf{s} = \sum_i \mathbf{f}_i \otimes \mathbf{r}_i \quad (11)$$

A dependency tree, for example, is naturally represented this way, with the roles being predicates paired with grammatical relations, such as *object of eat*, and the fillers being heads of arguments, such as *hotdog*. The vector representation of *eat hotdog* is  $\overrightarrow{\text{hotdog}} \otimes \overrightarrow{\text{eat\_obj}}$ . The vector representation of a whole dependency tree is the sum of the tensor products over all the dependency edges in the tree.

Smolensky and Legendre (2006) argue at length for why the tensor product is appropriate for combining connectionist and symbolic structures. A number of other proposals for realizing symbolic structures in vector representations are described, including

Plate (2000) and Pollack (1990). Smolensky’s claim is that, despite perhaps appearances to the contrary, all are special cases of a generalized tensor product representation.

Clark and Pulman (2007) suggest that a similar scheme to Smolensky’s could be applied to parse trees for NLP applications, perhaps with the use of a convolution kernel (Haussler 1999) to calculate similarity between two parse trees with distributed representations at each node. A similar idea has been fleshed out in Zanzotto, Ferrone, and Baroni (2015), showing how kernel functions can be defined for a variety of composition operations (including most of the operations discussed in this article), allowing the similarity between two parse trees to be computed based on the distributional representations at the leaves and the composed representations on the internal nodes.

There are similarities between Smolensky’s approach and the Categorical framework from Section 4.2, in that tensor representations are central to both. However, in Smolensky’s approach, the tensor product is the operation used for composition of the predicate and its argument, whereas tensor contraction is the composition operation in the Categorical framework.

## 5. Composition Methods – Implementation

This section gives a more detailed description of the methods we have tested on RELPRON, ranging from some trivial lexical baselines, through to the simple arithmetic operators, and finally more sophisticated methods based on the Categorical framework. Both count-based vectors and neural embeddings have been investigated.

### 5.1 Word Vectors

In this section we describe how we built the word and holistic phrase vectors used by the various composition methods.

*5.1.1 Count-based Vectors (Count).* For historical completeness we include classical count-based vectors in this article, although these large vectors are impractical for more sophisticated tensor learning. These vectors are used only for the baseline results (Section 6, Table 5), including the Frobenius algebra composition method of Sadrzadeh, Clark, and Coecke (2013).

Count-based word vectors were built from a 2013 download of Wikipedia, using the settings from Grefenstette and Sadrzadeh (2011). The top 2,000 words in the corpus (excluding stopwords) were used as features, with a co-occurrence window three content words either side of the target word (i.e. with stopwords removed before applying the window), and co-occurrence counts weighted as the probability of a context word given the target word divided by the overall probability of the context word (Mitchell and Lapata 2008). We found that these settings achieved the best result on the transitive verb composition task of Grefenstette and Sadrzadeh (2011).

*5.1.2 Reduced Count-based Vectors (Count-SVD).* As a variant on the classical count-based vectors, we also present baseline results (Section 6, Table 5) on a count-based space reduced with Singular Value Decomposition (SVD), which produces dense vectors more comparable to neural embeddings.

Word vectors were built from the same 2013 Wikipedia download, using the settings of Polajnar, Fagarasan, and Clark (2014). The top 10,000 words in the corpus (excluding stopwords) were used as features, with sentence boundaries defining the co-occurrence window, and co-occurrence counts weighted with the t-test statistic (Curran 2004). SVD

was used to reduce the number of dimensions to 300. Following Polajnar and Clark (2014), context selection (with  $N = 140$ ) and row normalization were performed prior to SVD.

*5.1.3 Neural Embeddings (Skip-Gram).* Our main results in Section 6 are produced using neural embeddings. These types of vectors have proven useful for a wide variety of tasks in recent literature, and in our early experiments with tensor learning provided the best results on the RELPRON development data and other composition tasks. For our methods, in addition to the word vectors, we required holistic vectors to use as targets for training the verb and relative pronoun matrices and tensors. The holistic vectors were also obtained using neural embeddings.

Word vectors were built using a re-implementation of skip-gram with negative sampling (Mikolov et al. 2013) on a lemmatized 2015 download of Wikipedia. We used a window of ten words on either side of the target, with ten negative samples per word occurrence, and 100-dimensional target and context vectors. All lemmas occurring fewer than 100 times in the corpus were ignored.<sup>13</sup> The context vectors produced during word vector learning were retained for use during training of the holistic vectors for phrases.

Verb matrices and relative pronoun matrices and tensors require holistic vectors to serve as training targets for linear regression. These vectors represent the observed contexts in which instances of verbs or relative pronouns with their arguments are found in the Wikipedia corpus. We tagged and parsed the corpus with the Stanford Parser (Chen and Manning 2014). For learning of verb matrices, we extracted verb-subject and verb-object pairs, ensuring that each pair occurred within a transitive verb construction. For learning of relative pronoun matrices and tensors, we extracted ⟨head noun, verb, argument⟩ tuples that occurred within relative clause constructions. To identify relative clauses, we looked for characteristic dependency subtrees containing the Stanford Universal Dependency labels RELCL and PRONTYPE=REL. The grammatical role of the relative pronoun argument of the verb, and of the verb's in situ argument, determined whether it was a subject or object relative clause. We used a small number of manually defined heuristics to help limit the results to true relative clauses, e.g. to exclude cases of clausal complementation on the verb, as in *an incident that left one government embarrassed*.

The contexts of our extracted verb-argument pairs and relative clause tuples were used to train holistic vectors using our implementation of skip-gram, only retaining phrases that occurred at least twice in the corpus. For holistic vectors, we followed Paperno, Pham, and Baroni (2014) and used a wider window of 15 words either side of the target. Skip-gram has a random component in the initialisation of vectors. In order to ensure that the learned holistic and word vectors were in the same vector space, we trained both using the same context vectors.

## 5.2 Verb Matrices

In this section we describe two methods used for constructing verb matrices for the composition methods based on the Categorical framework.

---

<sup>13</sup> The word *slipping* which occurs once in the RELPRON test set was missing from the resulting corpus. We substituted the vector for *slip* at test time.

5.2.1 *Relational Verb Matrices.* Using the relational method of Grefenstette and Sadrzadeh (2011), as described in Section 4.2.3, verb matrices were built as a sum of outer products of their subject-object (SO) pairs. Pairs were subject to a minimum per-verb count of two and a minimum frequency of 100 for both nouns in the 2013 Wikipedia download, and not weighted by frequency.

5.2.2 *Decoupled Verb Matrices with Linear Regression.* All of the Categorical methods require a representation for the transitive verb in the relative clause. Based on the CCG category  $(S \setminus NP) / NP$  of a transitive verb, this is a third-order tensor,  $S \otimes N \otimes N$ . The relative pronoun has a CCG type of  $(NP \setminus NP) / (S \setminus NP)$  or  $(NP \setminus NP) / (S / NP)$ , resulting in a fourth-order tensor,  $N \otimes N \otimes S \otimes N$ . In order to make the learning of a relative pronoun tensor tractable, we make use of the method introduced by Paperno, Pham, and Baroni (2014) and Polajnar, Fagarasan, and Clark (2014), in order to reduce the semantic type of the transitive verb to  $S \otimes N$ . This method models a transitive verb as two second-order tensors (matrices), one governing the interaction of the verb with its subject and the other with its object, and has been shown to reproduce sentence semantics without loss of accuracy. This “decoupling” approximation for the transitive verb allows the relative pronoun to be modelled as a third-order tensor,  $N \otimes N \otimes S$ .

Specifically, we adopt the approach of Paperno, Pham, and Baroni (2014). For each verb  $V$  we learn a pair of matrices  $\bar{V}^S, \bar{V}^O$ . These matrices are used by Paperno, Pham, and Baroni (2014) to calculate the vector for an SVO triple  $(s, V, o)$  as in Equation 12, which also gives a lexicalized example.

$$\bar{V}^S \vec{s} + \bar{V}^O \vec{o} + \vec{v} \quad \text{ex. } \overrightarrow{\text{eat}}^S \text{cat} + \overrightarrow{\text{eat}}^O \text{mouse} + \overrightarrow{\text{eat}} \quad (12)$$

We also adopt the modification of Gupta, Utt, and Padó (2015), in which the word vector for the verb is not added to the composition result, since this formulation was found to yield more accurate results. With this modification, we calculate the vector for an SVO triple  $(s, V, o)$  as in Equation 13.

$$\bar{V}^S \vec{s} + \bar{V}^O \vec{o} \quad \text{ex. } \overrightarrow{\text{eat}}^S \text{cat} + \overrightarrow{\text{eat}}^O \text{mouse} \quad (13)$$

Verb matrices were learned using  $L_2$  ridge regression (Hoerl and Kennard 1970), also known as  $L_2$ -regularised regression. For a given verb  $V$ , the training data for  $\bar{V}^S$  consisted of (subject vector, holistic subject-verb vector) pairs, and the training data for  $\bar{V}^O$  consisted of (object vector, holistic verb-object vector) pairs. In addition, each pair was weighted by the logarithm of its number of occurrences in the corpus, which we found improved performance on the RELPRON development set. Regression was implemented in python using standard NumPy operations, and we optimised the regularisation parameter on the RELPRON development set, using MAP scores and the SPLF composition method described in Section 5.5.3. A single regularisation parameter was used globally for all verbs and set to 75 after tuning.

### 5.3 Simple Composition Methods

In this section we describe our relative clause composition methods. We first present the baseline methods, consisting of lexical comparisons, simple arithmetic operators, and the Frobenius algebra method of Clark, Coecke, and Sadrzadeh (2013) and Sadrzadeh,

Clark, and Coecke (2013). We then introduce our main methods, all of which, like the Frobenius algebra method, are based on the Categorical framework, but which differ according to whether and how the relative pronoun is modelled. All of our main methods are implemented using neural embedding vectors.

**5.3.1 Lexical Baselines.** Two simple lexical similarity baselines (**Lexical**) set the vector representation of the property equal to the verb vector or the argument vector. No composition is involved. These baselines check whether a system can predict the similarity between, for example, *traveler* and *person that hotel serves* by the similarity between *traveler* and *hotel* (argument) or *traveler* and *serve* (verb). We do not have a lexical baseline consisting of the head noun, since this would not produce a well-defined ranking of properties, given that many properties share the same head noun (*person* in this example).

**5.3.2 Arithmetic Operators.** Arithmetic composition methods (**Arithmetic**) are the vector addition and elementwise vector product (Mitchell and Lapata 2008, 2010) of the vectors for the three lexical items in the property: head noun, verb, and argument. For vector addition, we also perform ablation to determine the contribution of each lexical item. The relative pronoun itself is not modelled in the composed representation for the arithmetic operators. Although the arithmetic operators are simple, vector addition has proven a difficult baseline to beat in many composition tasks.

## 5.4 Frobenius Algebra Baseline

The Frobenius algebra method of Clark, Coecke, and Sadrzadeh (2013) and Sadrzadeh, Clark, and Coecke (2013) is a previously existing implementation of the Categorical framework for relative clauses, and was described in detail in Section 4.3. The composition of the relative clause vector is shown in Equations 14 - 15, for subject and object relative clauses respectively, where  $\odot$  represents elementwise multiplication; note that unlike the other Categorical methods, there is a single verb matrix, which is learned as a relational matrix.


$$\vec{n} \odot (\overline{V} \vec{\sigma}) \quad \text{ex. } \overrightarrow{\text{device}} \odot (\overrightarrow{\text{detect}} \overrightarrow{\text{planet}}) \quad (\text{subject}) \quad (14)$$

$$\vec{n} \odot (\overline{V}^T \vec{s}) \quad \text{ex. } \overrightarrow{\text{device}} \odot (\overrightarrow{\text{play}}^T \overrightarrow{\text{shepherd}}) \quad (\text{object}) \quad (15)$$

## 5.5 Composition Methods within Categorical Framework

In the next sections we describe the main composition methods evaluated in this article. These methods are based on the Categorical framework, as is the Frobenius algebra baseline; however, the methods described here are all implemented with Skip-Gram vectors, and matrices and tensors learned by linear regression. Figure 3 gives a schematic view of these composition methods.

**5.5.1 Tensor Composition Method (RPTensor).** The **RPTensor** method models the relative pronoun as a third-order tensor  $\overline{\overline{R}}$ . Learning a relative pronoun tensor from holistic vectors for relative clauses was suggested by Baroni, Bernardi, and Zamparelli (2014), but to our knowledge this is the first implementation of a function word tensor. We learn

**Figure 3**

Schematic view of the major composition methods investigated in this article: (a) Addition, (b) RPTensor, (c) PLF, (d) SPLF, (e) FPLF. Figures are agnostic as to subject or object grammatical function for the relative clause.

two separate tensors  $\overline{\overline{R}}^S$  and  $\overline{\overline{R}}^O$  for subject and object relative clauses, respectively, since we assume that the relative pronoun may have a different semantics in each case.

The relative pronoun tensors have the type  $N \otimes N \otimes S$ , and combine via tensor contraction with the head noun vector and the vector resulting from verb-argument composition. Here and in subsequent sections,  $\vec{n}$  is the head noun vector, and  $\vec{a}$  is the argument vector, which may be the object (in a subject relative clause) or the subject (in an object relative clause). The composition of the relative clause vector is shown in Equations 16 - 17, for subject and object relative clauses respectively.

$$\vec{n} \overline{\overline{R}}^S (\overline{V}^O \vec{a}) \quad \text{ex. } \vec{\text{device}} \overline{\overline{\text{that}}}^S (\overline{\text{detect}}^O \vec{\text{planet}}) \quad (\text{subject}) \quad (16)$$

$$\vec{n} \overline{\overline{R}}^O (\overline{V}^S \vec{a}) \quad \text{ex. } \vec{\text{device}} \overline{\overline{\text{that}}}^O (\overline{\text{play}}^S \vec{\text{shepherd}}) \quad (\text{object}) \quad (17)$$

Note that the appropriate verb matrix, subject or object, is first contracted with its argument, so that the relative pronoun is composing the head noun with the verb-argument phrase (a vector).

As with the verb matrices, the relative pronoun tensors  $\overline{\overline{R}}^S$  and  $\overline{\overline{R}}^O$  are learned by weighted  $L_2$  ridge regression, where the training data consists of (head noun vector,

holistic verb-argument vector, holistic relative clause vector) tuples. The regularisation parameter was optimised on the RELPRON development set using MAP scores and set to 80 for both tensors.

The power of the tensor comes from the fact that it captures the interaction between the head noun and the composed verb-argument phrase, and hence models the relative pronoun in a meaningful way that composition methods such as vector addition, which omits the pronoun entirely, do not. However, capturing these interactions leads to a large number of parameters, and the sparsity of training data is therefore a challenge in training the full tensor model. Our corpus of 72M sentences contained only around 800K occurrences of relative clauses, compared to over 80M verb-subject and verb-object occurrences (4M types); and this was heavily unbalanced across subject and object relative clauses, with 771K subject relative clause occurrences (222K types) compared to 21K object relative clause occurrences (only 8K types).

**5.5.2 PLF Composition Method (PLF).** The **PLF** method implements the Practical Lexical Function model of Paperno, Pham, and Baroni (2014). As described in Section 5.2.2, this is one of the models which decouples the subject and object interaction of the transitive verb. In PLF, function words such as relative pronouns are explicitly treated as empty elements. A noun phrase containing a relative clause is not modelled as a noun phrase, but rather an SVO sentence. For example, *device that detects planets* is equivalent to *device detects planets*, and *device that shepherd plays* is equivalent to *shepherd plays device*. The subject and object verb matrices combine with their arguments by tensor contraction (in this case, matrix multiplication), and the resulting vectors are added. The composition of the relative clause vector is shown in Equations 18 - 19, for subject and object relative clauses respectively.

$$\overline{V}^S \overline{n} + \overline{V}^O \overline{a} \quad \text{ex. } \overline{\text{detect}}^S \overline{\text{device}} + \overline{\text{detect}}^O \overline{\text{planet}} \quad (\text{subject}) \quad (18)$$

$$\overline{V}^S \overline{a} + \overline{V}^O \overline{n} \quad \text{ex. } \overline{\text{play}}^S \overline{\text{shepherd}} + \overline{\text{play}}^O \overline{\text{device}} \quad (\text{object}) \quad (19)$$

As noted in Section 5.2.2, in this and all other PLF variants we use the modification of Gupta, Utt, and Padó (2015) to the original PLF, so the word vector for the verb is not included in the composition.

**5.5.3 Simplified PLF Composition Method (SPLF).** The **SPLF** method is a simplification of PLF, where the verb and argument are combined by tensor contraction as in PLF, but the resulting verb-argument vector is combined with the head noun by vector addition. The relative pronoun is not explicitly modelled, although we might think of the vector addition as a crude representation of relative pronoun semantics that “distinguishes” the roles played by the extracted head noun and the in situ argument. SPLF is also similar to Addition, except that the verb and argument are combined by tensor contraction (in this case, matrix multiplication). The composition of the relative clause vector is shown in Equations 20 - 21, for subject and object relative clauses respectively.

$$\overline{n} + \overline{V}^O \overline{a} \quad \text{ex. } \overline{\text{device}} + \overline{\text{detect}}^O \overline{\text{planet}} \quad (\text{subject}) \quad (20)$$

$$\overline{n} + \overline{V}^S \overline{a} \quad \text{ex. } \overline{\text{device}} + \overline{\text{play}}^S \overline{\text{shepherd}} \quad (\text{object}) \quad (21)$$

**5.5.4 Full PLF Composition Method (FPLF).** The **FPLF** method offers a way to mitigate the sparsity of the training data for RPTensor, while retaining explicit modelling of the relative pronoun. Here we extend the intuition about modelling higher order tensors with matrices to the relative pronoun itself, by decoupling the interaction of the pronoun with the head noun from its interaction with the composed verb-argument phrase. We had already reduced the relative pronoun to a third-order tensor, and we now represent it by two matrices, one for head noun interaction and one for verb-argument interaction. We further separate subject and object relative clauses, and thus end up with two pairs of matrices  $(\overline{R}_n^S, \overline{R}_{va}^S)$  and  $(\overline{R}_n^O, \overline{R}_{va}^O)$ . We call this a “full” variant of PLF because it extends the PLF intuition to function words, a suggestion made by Paperno, Pham, and Baroni (2014) but which to our knowledge has not been previously implemented. The composition of the relative clause vector is shown in Equations 22 - 23, for subject and object relative clauses respectively.

$$\overline{R}_n^S \vec{n} + \overline{R}_{va}^S (\overline{V}^O \vec{a}) \quad \text{ex. } \overline{\text{that}}_n^S \overline{\text{device}} + \overline{\text{that}}_{va}^S (\overline{\text{detect}}^O \overline{\text{planet}}) \quad (\text{subject}) \quad (22)$$

$$\overline{R}_n^O \vec{n} + \overline{R}_{va}^O (\overline{V}^S \vec{a}) \quad \text{ex. } \overline{\text{that}}_n^O \overline{\text{device}} + \overline{\text{that}}_{va}^O (\overline{\text{play}}^S \overline{\text{shepherd}}) \quad (\text{object}) \quad (23)$$

The relative pronoun matrices  $(\overline{R}_n^S, \overline{R}_{va}^S)$  and  $(\overline{R}_n^O, \overline{R}_{va}^O)$  are learned by weighted  $L_2$  ridge regression, where the training data consists of (noun vector, holistic relative clause vector) pairs for  $(\overline{R}_n^S, \overline{R}_n^O)$ , and (holistic verb-argument vector, holistic relative clause vector) pairs for  $(\overline{R}_{va}^S, \overline{R}_{va}^O)$ . The regularisation parameter was optimised on the RELPRON development set and set to 75 for all four matrices.

**5.5.5 Categorical Baselines.** We also investigate two Categorical baselines. **VArg** is the verb matrix (subject or object, as appropriate) multiplied by the in situ argument, without the head noun. For example, *person that rescuer assists* is *rescuer assists*, while *device that detects planets* is *detects planets*. The composition of the relative clause vector is shown in Equations 24 - 25, for subject and object relative clauses respectively.

$$\overline{V}^O \vec{a} \quad \text{ex. } \overline{\text{detect}}^O \overline{\text{planet}} \quad (\text{subject}) \quad (24)$$

$$\overline{V}^S \vec{a} \quad \text{ex. } \overline{\text{play}}^S \overline{\text{shepherd}} \quad (\text{object}) \quad (25)$$

**Vhn** is the verb matrix (subject or object, as appropriate) multiplied by the head noun. In this case, *person that rescuer assists* is *assists person*, while *device that detects planets* is *device detects*. The composition of the relative clause vector is shown in Equations 26 - 27, for subject and object relative clauses respectively.

$$\overline{V}^S \vec{n} \quad \text{ex. } \overline{\text{detect}}^S \overline{\text{device}} \quad (\text{subject}) \quad (26)$$

$$\overline{V}^O \vec{n} \quad \text{ex. } \overline{\text{play}}^O \overline{\text{device}} \quad (\text{object}) \quad (27)$$

## 6. Evaluation and Results

This section describes the evaluation of the ranking task on RELPRON and presents the results of the various composition methods.

**Table 5**

MAP scores of Lexical, Arithmetic, and Frobenius algebra methods on the RELPRON development set using Count, Count-SVD, and Skip-Gram vectors, and relational verb matrices.

	Method	Count	Count-SVD	Skip-Gram
<b>Lexical</b>	$\vec{arg}$	0.272	0.386	0.347
	$\vec{verb}$	0.138	0.179	0.176
<b>Arithmetic</b>	$\vec{hn} \odot \vec{arg} \odot \vec{verb}$	0.364	0.123	0.181
	$\vec{hn} + \vec{arg} + \vec{verb}$	<b>0.386</b>	<b>0.442</b>	<b>0.496</b>
	$\vec{arg} + \vec{verb}$	0.331	0.407	0.401
	$\vec{hn} + \vec{arg}$	0.339	0.425	0.450
	$\vec{hn} + \vec{verb}$	0.231	0.229	0.264
<b>Categorical</b>	Frobenius	0.277	0.023	0.030

## 6.1 Evaluation

To evaluate a composition method, composed vectors are produced for all properties in the dataset. For each term, the properties are ranked according to cosine similarity with the term vector. Evaluation is based on Mean Average Precision (MAP), which is defined as:

$$\text{MAP} = \frac{1}{N} \sum_{i=1}^N AP(t_i) \quad (28)$$

where  $N$  is the number of terms in the dataset, and  $AP(t)$  is the Average Precision for term  $t$ , defined as:

$$AP(t) = \frac{1}{P_t} \sum_{k=1}^M \text{Prec}(k) \times \text{rel}(k) \quad (29)$$

where  $P_t$  is the number of correct properties for term  $t$  in the dataset,  $M$  is the total number of properties in the dataset,  $\text{Prec}(k)$  is the precision at rank  $k$ , and  $\text{rel}(k)$  is an indicator function which is equal to one if the property at rank  $k$  is a correct property for  $t$ , and zero otherwise.

## 6.2 Results

Table 5 shows the results for the baseline methods, using Count, Count-SVD, and Skip-Gram vectors. The highest MAP score for all three types of vectors is achieved with vector addition, adding the vectors for head noun, verb, and argument, and the Skip-Gram vectors achieve the best result overall, with a MAP of 0.496. This is already a challenging score to beat, since it represents a system which on average places a correct property at every second position in the top-ranked properties for each term. Elementwise multiplication performs almost as well for Count vectors, but very poorly

**Table 6**

MAP scores of composition methods on the RELPRON development and test sets using Skip-Gram vectors. \* Significantly higher than next lowest result ( $p < 0.05$ ).

Method	Development	Test
SPLF	<b>0.496</b>	<b>0.497</b>
Addition	<b>0.496*</b>	<b>0.472*</b>
PLF	0.444	0.433*
FPLF	0.405	0.387*
RPTensor	0.380	0.354
VArg	0.448	0.397
Vhn	0.219	0.204

for Count-SVD and Skip-Gram vectors, which is consistent with previous findings on elementwise multiplication in dense vector spaces (Vecchi, Baroni, and Zamparelli 2011; Utsumi 2012; Polajnar and Clark 2014; Milajevs et al. 2014). We will therefore use vector addition as our primary arithmetic vector operation for the rest of the article. Moreover, we will use Skip-Gram vectors for our main Categorical framework-based experiments.

Turning to the lexical baselines and the ablation tests for vector addition, we may assess the relative contribution to the vector addition result of the three lexical items in the relative clause. We look first at the lexical baselines. Similarity between the term and argument achieves a MAP of 0.347 for Skip-Gram vectors, while similarity between the term and the verb produces a much worse ranking. Looking at the arithmetic operators, the ablation tests show that the head noun, verb, and argument all contribute to the correct ranking of properties. However, the argument is the most important, as the sum of head noun and verb vectors achieves a lower MAP (0.264 for Skip-Gram vectors) than the other two combinations (0.401 and 0.450, respectively). The importance of the argument is fairly intuitive; for example, in *family: organization that claims ancestry*, the words *family* and *ancestry* are clearly the most closely related. However, Table 5 shows that the argument alone is insufficient to rank properties and a method must perform some composition to achieve a respectable MAP score.

In Table 5 we also see that the Frobenius method, which uses a relational matrix for the verb, is not competitive with vector addition. It achieves a MAP of only 0.277 using Count vectors, and in the other two vector spaces, the MAP is almost zero, presumably due to the elementwise multiplication that is part of the method. Based on this result, we do not pursue the Frobenius algebra method further.

Table 6 shows the results of the composition methods and Categorical baselines on the RELPRON development and test data. The relative performance of the methods is very similar across the development and test portions of the dataset. In both cases, SPLF and Addition are the top two performers, with SPLF having a higher MAP than Addition on the test data, although the difference is not significant. Both methods achieve a MAP of nearly 0.50, which corresponds to putting a correct property in every second place on average.

The next highest MAP is achieved by PLF. It is instructive to compare PLF to the Categorical baselines, VArg and Vhn, since PLF sums these two components. Vhn performs very badly, in line with the sum of the verb and head noun vectors in Table 5. However, the head noun does contribute to the composition of relative clause meaning, as we can see from the fact that on the test data, PLF achieves a significantly higher MAP

than VArg. It is also relevant to compare PLF with SPLF. PLF learns a verb matrix for the interaction with the head noun, while SPLF approximates the head noun interaction as addition. SPLF significantly outperforms PLF on the development and test data, a difference which suggests that modelling the relative clause as an SVO sentence may not be the optimal approach for relative clause composition.

FPLF and RPTensor are the two methods which learn a relative pronoun function by regression. FPLF achieves a lower MAP than SPLF, Addition, and PLF, but a higher MAP than RPTensor, supporting the hypothesis that the available training data in the form of holistic vectors for relative clauses is insufficient for the number of parameters in the full tensor method. Both methods underperform the VArg functional baseline and some of the arithmetic baselines.

## 7. Error Analysis

In this section we examine in detail the performance of some of the main composition methods on the RELPRON development data. We are interested in four main points. First, we investigate whether the grammatical function (subject or object) of the relative clause plays a role in the MAP score. We find that most of the methods are relatively well-balanced across grammatical functions, but FPLF shows a lower MAP on object properties, where the amount of training data is substantially lower.

Second, we investigate whether the different head nouns in the dataset have any effect on MAP score. We find that the scores are relatively well-balanced across head nouns, but that more concrete terms and their properties may be easier to model. However, based on qualitative observations a more important factor seems to be term polysemy.

Third, we look at how well the composition methods capture the intersective semantics of the relative clause construction. We break down the evaluation into two independent components: how often the methods are able to identify the correct head noun, and how often they are able to pick out correct properties when the head noun is known. We find that all methods do very well at identifying the correct head noun, but FPLF lags behind the others at picking the correct property.

Finally, we look at some common errors on the RELPRON development data. One source of errors is lexical overlap in terms and properties. We also observe that some of the systems assign high ranks to plausible but unannotated properties, and address this issue with an evaluation scenario where the property is the query and the task is to rank the correct term highest. We find that on average all the methods rank the correct term at approximately position two.

We focus in this section on four methods: SPLF and Addition, the two methods with the highest overall MAP score; FPLF, the better-performing method of the two which explicitly learn a relative pronoun representation; and VArg, as a high-performing Categorical baseline.

### 7.1 Accuracy by Grammatical Function

An obvious first step in understanding the RELPRON results is to break them down by grammatical function. Since subjects and objects are asymmetric in their interaction with the verb — in particular, verbs show stronger selectional preferences for objects than for subjects (Marantz 1984; Kratzer 1996, 2002) — we might expect that relative clauses with different extraction sites are not equally easy to compose. In addition, for the methods that learn an explicit representation of the relative pronoun (FPLF in

**Table 7**

MAP scores by grammatical function of extracted element in relative clause, on development set.

	<b>SPLF</b>	<b>Addition</b>	<b>FPLF</b>	<b>VArg</b>
Subject	0.585	0.571	0.570	0.530
Object	0.541	0.574	0.428	0.489

this analysis), the substantial difference in the amount of training data might affect the result.

The MAP scores for each grammatical function are shown in Table 7. The first row shows the subject results, in which only the subject properties from the development data are ranked by similarity for each term. The Average Precision (AP) is calculated for each term, and MAP is calculated as the mean AP over over the set of all terms.<sup>14</sup> The object results are calculated analogously, although note that the MAP scores in Table 7 are not directly comparable to the scores on the full development set, because the number of confounders is lower for each term.

For SPLF and Addition, Table 7 shows that the MAP scores are well balanced. For FPLF, on the other hand, the MAP scores are unbalanced: 0.570 for subject properties, in line with SPLF and Addition; and only 0.428 for object properties. This difference provides support for the hypothesis that the amount of training data available for subject vs. object relative clauses makes a difference to the resulting accuracy.<sup>15</sup> The term *navy* provides an example of the differential performance between subject and object properties. SPLF ranks four correct subject and three correct object properties within its top ten ranked properties, while FPLF ranks five correct subject properties and only one correct object property within its top ten.<sup>16</sup>

The MAP scores for VArg are also unbalanced in favour of subject relative clauses, though the effect is not as pronounced. We attribute this difference to the strength of the selectional preference of a verb for its object compared to its subject, since verb-object combinations (for subject properties) may have greater semantic content, enabling a more accurate ranking of properties.

## 7.2 Accuracy by Head Noun

Breaking down the results on RELPRON by head noun – for example, *device* vs. *quality* – is another obvious place to look for differential performance in the composition methods. Table 8 shows the MAP scores by head noun on the development set. The head nouns are ordered from left to right according to the average concreteness rating of their terms, based on the ratings of Brysbaert, Warriner, and Kuperman (2014). In Table 8, the AP for

<sup>14</sup> Any term which has no properties with the relevant grammatical function is omitted from the calculation; for example, *accuracy* has no subject properties, so it is not included in the MAP for Subject.

<sup>15</sup> As further support for this hypothesis, the RPTensor method (not in the table) showed the same discrepancy, with a MAP of 0.520 on subject properties, and 0.400 on object properties.

<sup>16</sup> The correct properties for SPLF are, in order, *organization that uses submarine* (S), *organization that blockades port* (S), *organization that fleet destroys* (O), *organization that battleship fights* (O), *organization that maintains blockade* (S), *organization that vessel serves* (O), *organization that defeats fleet* (S). The correct properties for FPLF are, in order, *organization that uses submarine* (S), *organization that blockades port* (S), *organization that maintains blockade* (S), *organization that fleet destroys* (O), *organization that establishes blockade* (S), *organization that defeats fleet* (S).

**Table 8**

MAP scores by head noun on development set. AP is calculated over all properties for each term, and mean AP calculated for each head noun. Head nouns are ordered from left to right by increasing concreteness.

	Head Noun						
	quality	organization	person	document	building	player	device
<b>SPLF</b>	0.490	0.500	0.490	0.393	0.592	0.595	0.462
<b>Addition</b>	0.464	0.485	0.394	0.463	0.632	0.546	0.512
<b>FPLF</b>	0.279	0.471	0.415	0.360	0.502	0.448	0.378
<b>VArg</b>	0.398	0.563	0.413	0.390	0.542	0.443	0.384

each term is based on a ranking of all properties in the development set, and the MAP for a head noun is the mean AP over all its terms.<sup>17</sup>

Table 8 shows that the MAP scores are relatively well-balanced across head nouns for all methods, suggesting that there are no real outliers among head nouns in the ability of the methods to compose relative clauses. The general distribution of scores across head nouns for the four methods is also roughly similar. *Building* and *player* are consistently high, with *building* in particular having a MAP of at least 0.5 for all methods. These head nouns have some of the most concrete terms, which might suggest that concrete terms have properties which are easier to model, or at least that concrete terms have better-quality word vectors which makes the ranking easier. This is not a clear effect, however, since the methods mostly have average-to-high MAPs for *quality* and *organization*, which have terms with low concreteness, and average-to-low MAPs for *device*, which has the most concrete terms in the development set.

A factor which appears to play a role in MAP scores is term polysemy. All of the methods score low on *document*, and SPLF in particular achieves its lowest MAP for this head noun. Several document terms (*lease*, *form*, *assignment*, *bond*) exhibit very low AP. The term *form* provides an example; it has an extremely low AP of 0.01. Although *form* was chosen by the annotators in its *document* sense, with correct properties including *document that parent signs* and *document that applicant completes*, other senses may be more dominant in the source corpus, confusing the similarity ranking.<sup>18</sup> Prior disambiguation of words in the dataset, as in Kartsaklis and Sadrzadeh (2013), might improve performance.

We believe that polysemy may come into play with *device* terms as well, though the terms have high concreteness. One term with a very low AP is *watch*. Among the properties that SPLF ranks high for *watch* are *person that police hunt (killer)* and *device that views stars (telescope)*, both of which are related to different senses of *watch*.

### 7.3 Capturing the Semantics of Relative Clauses

A system that successfully captures the semantics of relative clauses must integrate the semantic contribution of the head noun with the contributions of the verb and argument; for example, identifying that a *saw* is a *device*, and that among devices, it

<sup>17</sup> Unlike in Section 7.1, these MAP scores are directly comparable to those on the full development set, since the number of confounders is the same.

<sup>18</sup> The top four properties for *form* as ranked by SPLF are: *organization that undergoes merger (division)*, *organization that siblings form (family)*, *organization that infantry reinforces (army)*, *organization that restructuring creates (division)*.

**Table 9**

Average proportion of top 10 ranked properties that have the correct head noun, on development set.

	Head Noun							Overall
	quality	organization	person	document	building	player	device	
<b>SPLF</b>	0.87	0.71	0.69	0.68	0.86	1.00	0.83	0.79
<b>Addition</b>	0.84	0.75	0.60	0.64	0.88	1.00	0.78	0.77
<b>FPLF</b>	0.84	0.71	0.62	0.72	0.88	0.98	0.91	0.79

**Table 10**

MAP scores within head nouns on development set. Only the properties with the correct head noun are ranked for each term.

	Head Noun						
	quality	organization	person	document	building	player	device
<b>SPLF</b>	0.638	0.738	0.616	0.511	0.666	0.596	0.564
<b>Addition</b>	0.587	0.686	0.563	0.602	0.711	0.551	0.580
<b>FPLF</b>	0.421	0.683	0.579	0.466	0.569	0.467	0.424

is the one that *cuts wood*. In this section we break down the results on the RELPRON development set into two measures that demonstrate performance on these subtasks independently.

We look first at how often the methods are able to identify the correct head noun. We consider the top ten ranked properties for each term from the full development set, and calculate the percentage of them which have the correct head noun, regardless of whether the whole property is correct. The results are shown in Table 9, where we omit the VArg method since it does not take the head noun into account. Table 9 shows that overall, nearly eight out of the top ten properties have the correct head noun. FPLF matches the performance of the other two methods, although it is the only one of the three that incorporates the head noun by an operation other than addition.

We next look at the MAP scores when the ranking of properties for each term is restricted to properties with the correct head noun. For example, given a *building* term such as *ruin*, the methods must rank only the *building* properties. The task here is to distinguish *ruin* from *cinema*, *pub*, *observatory*, *house*, *mosque*, etc. The results are shown in Table 10.<sup>19</sup> We again omit VArg, since SPLF and Addition are equivalent to VArg when the head noun is held fixed.

Table 10 shows that SPLF and Addition perform similarly on distinguishing properties of terms that share a head noun, while FPLF lags behind. Combined with Table 9, this result shows that FPLF’s overall lower MAP on the full development set is due to difficulty distinguishing terms from one another when the head noun is known, not to difficulty in identifying the correct head noun.

Comparing Tables 9 and 10 highlights cases where the methods struggle with one relative clause composition subtask or the other, an effect which appears to vary by

<sup>19</sup> One question that might be asked is whether the ranking task within head nouns would be more difficult than the full ranking task, since all buildings (for example) share certain characteristics which might make their properties particularly good mutual confounders. In practice, this is difficult to measure, since the MAP scores in Table 10 are higher than the MAPs on the full development set because there are fewer confounders.

head noun. For example, all methods demonstrate perfect or near-perfect ability to rank *player* properties at the top for *player* terms, but average-to-low ability to distinguish players from one another. This may be a feature of the topic domain, in that the activities undertaken by different kinds of sports players — *golfer, batter, pitcher, quarterback, etc.* — are distributionally similar. On the other hand, all methods exhibit only average ability to identify the correct head noun for *organization* terms, but relatively high ability to select the correct *organization* properties when the head noun is known. For example, the SPLF ranking for *religion* shows a large amount of confusion with *person* properties, such as *person that defends rationalism (philosopher)*.<sup>20</sup> However, when restricted to ranking *organization* properties, as in Table 10, SPLF achieves a perfect MAP of 1.0 for *religion*.

#### 7.4 Common Errors and Evaluation with Properties as Queries

In addition to poor performance on polysemous terms (Section 7.2), we observed two common sources of error on RELPRON. The first is lexical overlap between terms and properties, a type of confounder that we deliberately included in the dataset (see Section 3.5). For example, the top two ranked properties for *batter* for all three methods are *player that walks batter (pitcher)* and *player that strikes batter (pitcher)*; and the top ranked property for *balance* for all three methods is *document that has balance (account)*. FPLF, despite its more sophisticated modelling of the relative pronoun, suffers from lexical overlap about as much as SPLF and Addition. However, we note that the RPTensor method is able to overcome this problem to a small extent, with the lexically confounding properties ranked a few positions below the top.

The second source of error is that the methods often assign high similarity to properties which are plausible descriptions for a term, but not annotated as gold positives. For example, the top-ranked property for *lease* for the RPTensor method is *document that government sells (bond)*. The highest-ranked properties for *intellectual* for SPLF include *person that promotes idea (advocate)* and *person that analyzes ontology (philosopher)*.<sup>21</sup> This phenomenon speaks to the difficulty of collecting distinctive properties during the manual annotation phase; future work may involve re-annotation of the dataset for additional positives. At present, the effect is to artificially lower the MAP ceiling, which affects all methods equally.

To address this issue, we performed another type of evaluation, this time treating properties as queries and ranking terms by their similarity to a property. We used Mean Reciprocal Rank (MRR) as the evaluation measure, since there is only one correct term per property. MRR is given by the formula:

$$\text{MRR} = \frac{1}{M} \sum_{p=1}^M \frac{1}{\text{rank}(p)} \quad (30)$$

where  $M$  is the number of properties in the dataset and  $\text{rank}(p)$  is the rank at which the correct term  $t$  is found for property  $p$ . The reasoning behind the MRR evaluation is

<sup>20</sup> The top five properties in the SPLF ranking are *person that defends rationalism (philosopher)*, *person that religion has (follower)*, *person that questions theology (philosopher)*, *person that teaches epistemology (philosopher)*, and *person that accepts philosophy (follower)*.

<sup>21</sup> Some of the highly-ranked but incorrect properties are surprisingly plausible, while clearly not being identifying properties for the term in question: SPLF ranks *building that sells popcorn (theater)* in second position for *pub*.

**Table 11**

Mean Reciprocal Rank (MRR) of correct term when property is used as query, on development set.

	Head Noun							Overall
	quality	organization	person	document	building	player	device	
<b>SPLF</b>	0.531	0.615	0.559	0.525	0.535	0.506	0.458	0.539
<b>Addition</b>	0.589	0.591	0.461	0.613	0.692	0.535	0.575	0.579
<b>FPLF</b>	0.462	0.498	0.457	0.434	0.518	0.427	0.409	0.461
<b>VArg</b>	0.606	0.662	0.547	0.567	0.619	0.556	0.545	0.591

that even if a property is applicable to more than one term, there is always a single term which is *most* similar to it. The results are shown in Table 11. For all three methods the correct term appears on average at approximately rank two. The MRR evaluation is also another use of RELPRON which can be exploited in future work.

## 8. Conclusions

RELPRON was designed to test compositional distributional semantic methods on a phrase type of intermediate size and complexity: more complex than two-word combinations, and involving a function word, but less complex than full sentences. The dataset is challenging, but results are promising with current techniques: a MAP score of 0.5 corresponds to every other property being ranked correctly, given a term. Again, in line with a growing body of literature, we observe that vector addition performs extremely well, but we are able to match its performance with a more complex method based on the Practical Lexical Function model.

Based on the qualitative analysis, our contention is that, in order to improve substantially on the results presented in this article, methods more linguistically sophisticated than vector addition will be required. We base this contention on a few factors. First, vector addition has little room for improvement, with potential gains limited to those that come from improving the overall quality of the vectors. It seems unlikely that vector addition can achieve a MAP much higher than its current one of around 0.5. We also note that the relative clauses used in this article are well below the ten-word length at which the quality of composed representations with vector addition has been suggested to degrade, so future datasets focusing on more complex linguistic constructions may require alternative methods.

For the more complex methods, on the other hand, there are likely to be readily achievable gains on RELPRON from increasing or changing the training data. This article represents the first large-scale implementation of methods that learn explicit Categorical representations for relative pronouns, but the training data in the form of holistic vectors for relative clauses was still observed to be relatively small compared to the number of parameters learned by the methods, especially for object relative clauses. Larger corpora or different sources of data — for example, including dictionary definitions or paraphrase datasets in the training data — may improve the models.

Beyond the issue of training data, the qualitative analysis showed that the various methods were sometimes unable to integrate the semantic contributions of the head noun and the remainder of the relative clause. Vector addition will not be able to capture this relationship any better than it already does, but other learning methods, for example involving non-linear models, may be able to do so.

It is also possible that a learning method designed more specifically for the RELPRON ranking task would perform better, especially since many of the training examples in our existing data are non-restrictive relative clauses; although this would require an alternative data source, and RELPRON itself is too small to provide training data. We expect RELPRON and similar datasets to be important evaluation tools for future methods that combine formal and distributional semantics, and hope that the insights provided by RELPRON inspire new work focused on linguistically challenging grammatical constructions.

### Acknowledgments

Laura Rimell and Stephen Clark were supported by EPSRC grant EP/I037512/1. Jean Maillard is supported by an EPSRC Doctoral Training Grant and a St John's Scholarship. Laura Rimell, Tamara Polajnar and Stephen Clark are supported by ERC Starting Grant DisCoTex (306920).

### References

- Agirre, Eneko. 2015. SemEval-2015 Task 2: Semantic Textual Similarity.
- Agirre, Eneko, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. SemEval-2012 Task 6: A pilot on Semantic Textual Similarity. In *Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval)*, Montreal.
- Agirre, Eneko, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and WeiWei Guo. 2013. \*SEM 2013 Shared Task: Semantic Textual Similarity. In *The Second Joint Conference on Lexical and Computational Semantics (\*SEM)*, Atlanta, GA.
- Baroni, M., G. Dinu, and G. Kruszewski. 2014. Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 238–247.
- Baroni, M. and R. Zamparelli. 2010. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Proceedings of the EMNLP Conference*, pages 1183–1193, Boston, MA.
- Baroni, Marco, Raffaella Bernardi, and Roberto Zamparelli. 2014. Frege in space: A program for compositional distributional semantics. *Linguistic Issues in Language Technology*, 9:5–110.
- Bender, Emily M., Dan Flickinger, Stephan Oepen, and Yi Zhang. 2011. Parser evaluation over local and non-local deep dependencies in a large corpus. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP 2011)*, pages 397–408.
- Bernardi, R., G. Dinu, M. Marelli, and M. Baroni. 2013. A relatedness benchmark to test the role of determiners in compositional distributional semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL): Short Papers*, pages 53–57.
- Biemann, C. and E. Giesbrecht. 2011. Distributional semantics and compositionality 2011: Shared task description and results. In *Proceedings of the Workshop on Distributional Semantics and Compositionality at ACL 2011*.
- Boleda, G., M. Baroni, L. McNally, and N. Pham. 2013. Intensionality was only alleged: On adjective-noun composition in distributional semantics. In *Proceedings of IWCS (International Conference on Computational Semantics)*, pages 35–46.
- Boleda, G., E. M. Vecchi, M. Cornudella, and L. McNally. 2012. First order vs. higher order modification in distributional semantics. In *Proceedings of EMNLP-CoNLL*, pages 1223–1233, Jeju Island, Korea.
- Brysbaert, Marc, Amy Beth Warriner, and Victor Kuperman. 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior research methods*, pages 904–911.
- Burnard, Lou. 2007. The British National Corpus, version 3. Distributed by Oxford University Computing Services on behalf of the BNC Consortium.
- Chen, Danqi and Christopher D Manning. 2014. A fast and accurate dependency parser using neural networks. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Cheung, Jackie C.K. and Gerald Penn. 2012. Evaluating distributional models of semantics for syntactically invariant inference. In *Proceedings of the 13th Meeting of the EACL*, Avignon, France.

- Clark, Stephen. 2013. Type-driven syntax and semantics for composing meaning vectors. In Chris Heunen, Mehrnoosh Sadrzadeh, and Edward Grefenstette, editors, *Quantum Physics and Linguistics: A Compositional, Diagrammatic Discourse*. Oxford University Press.
- Clark, Stephen. 2015. Vector space models of lexical meaning. In Shalom Lappin and Chris Fox, editors, *Handbook of Contemporary Semantic Theory - second edition*. Wiley-Blackwell, chapter 16.
- Clark, Stephen, Bob Coecke, and Mehrnoosh Sadrzadeh. 2008. A compositional distributional model of meaning. In *Proceedings of QI-2008*, pages 133–140, Oxford, UK.
- Clark, Stephen, Bob Coecke, and Mehrnoosh Sadrzadeh. 2013. The Frobenius anatomy of relative pronouns. In *Proceedings of the 13th Meeting on Mathematics of Language (MoL), special interest group of ACL*.
- Clark, Stephen and James R. Curran. 2007. Wide-coverage efficient statistical parsing with CCG and log-linear models. *Computational Linguistics*, 33(4):493–552.
- Clark, Stephen and Stephen Pulman. 2007. Combining symbolic and distributional models of meaning. In *Proceedings of AAAI Spring Symposium on Quantum Interaction*, Stanford, CA. AAAI Press.
- Coecke, Bob, Mehrnoosh Sadrzadeh, and Stephen Clark. 2010. Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis 36: A Festschrift for Joachim (Jim) Lambek*.
- Curran, James R. 2004. *From Distributional to Semantic Similarity*. Ph.D. thesis, University of Edinburgh.
- Dowty, D.R., R.E. Wall, and S. Peters. 1981. *Introduction to Montague Semantics*. Dordrecht.
- Fodor, Jerry and Zenon Pylyshyn. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28:3–71.
- Fried, Daniel, Tamara Polajnar, and Stephen Clark. 2015. Low-rank tensors for verbs in compositional distributional semantics. In *Proceedings of the Short Papers of ACL 2015*, pages 731–736, Beijing, China.
- Garrette, Dan, Katrin Erk, and Raymond Mooney. 2011. Integrating logical representations with probabilistic information using Markov Logic. In *Proceedings of IWCS*, Oxford, UK.
- Grefenstette, Edward. 2013. *Category-Theoretic Quantitative Compositional Models of Natural Language Semantics*. Ph.D. thesis, University of Oxford.
- Grefenstette, Edward, Georgiana Dinu, Yao-Zhong Zhang, Mehrnoosh Sadrzadeh, and Marco Baroni. 2013. Multi-step regression learning for compositional distributional semantics. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS-13)*, Potsdam, Germany.
- Grefenstette, Edward and Mehrnoosh Sadrzadeh. 2011. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1394–1404, Edinburgh, Scotland, UK, July.
- Grefenstette, Edward, Mehrnoosh Sadrzadeh, Stephen Clark, Bob Coecke, and Stephen Pulman. 2011. Concrete sentence spaces for compositional distributional models of meaning. In *Proceedings of the 9th International Conference on Computational Semantics (IWCS-11)*, pages 125–134, Oxford, UK.
- Gupta, Abhijeet, Jason Utt, and Sebastian Padó. 2015. Dissecting the practical lexical function model for compositional distributional semantics. In *Proceedings of STARSEM*, pages 153–158, Denver, CO.
- Haussler, David. 1999. Convolution kernels on discrete structures. Technical Report UCSC-CRL-99-10, Computer Science Department, University of California in Santa Cruz.
- Herbelot, Aurelie and Ann Copestake. 2015. Lexicalised compositionality. University of Cambridge Computer Laboratory, unpublished manuscript.
- Hermann, Karl Moritz, Edward Grefenstette, and Phil Blunsom. 2013. “Not not bad” is not “bad”: A distributional account of negation. In *Proceedings of the ACL Workshop on Continuous Vector Space Models and their Compositionality*.
- Hockenmaier, Julia and Mark Steedman. 2007. CCGbank: a corpus of CCG derivations and dependency structures extracted from the Penn Treebank. *Computational Linguistics*, 33(3):355–396.
- Hoerl, A.E. and R.W. Kennard. 1970. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67.
- Jurgens, David, Mohammad Taher Pilehvar, and Roberto Navigli. 2014. SemEval-2014 Task 3: Cross-level semantic similarity.

- Kartsaklis, Dimitri and Mehrnoosh Sadrzadeh. 2013. Prior disambiguation of word tensors for constructing sentence vectors. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Seattle, USA, October.
- Kartsaklis, Dimitri and Mehrnoosh Sadrzadeh. 2014. A study of entanglement in a categorical framework of natural language. In *Proceedings of the 11th Workshop on Quantum Physics and Logic (QPL)*, Kyoto, Japan, June.
- Kartsaklis, Dimitri, Mehrnoosh Sadrzadeh, and Stephen Pulman. 2012. A unified sentence space for categorical distributional-compositional semantics: Theory and experiments. In *Proceedings of COLING*, pages 549–558.
- Kiros, Ryan, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. Skip-thought vectors. *arXiv:1506.06726 [cs.CL]*.
- Kolda, T. G. and B. W. Bader. 2009. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500.
- Kratzer, A. 1996. Severing the external argument from its verb. In *Phrase structure and the lexicon*. Springer, Netherlands, pages 109–137.
- Kratzer, A. 2002. The event argument and the semantics of verbs. Ms., University of Massachusetts, Amherst.
- Krishnamurthy, Jayant and Tom M. Mitchell. 2013. Vector space semantic parsing: A framework for compositional vector space models. In *Proceedings of the 2013 ACL Workshop on Continuous Vector Space Models and their Compositionality*, Sofia, Bulgaria.
- Lambek, Joachim. 2008. *From Word to Sentence. A Computational Algebraic Approach to Grammar*. Polimetrica.
- Lazaridou, A., E. Vecchi, and M. Baroni. 2013. Fish transporters and miracle homes: How compositional distributional semantics can help NP parsing. In *Proceedings of EMNLP*, pages 1908–1913.
- Levy, Omer, Yoav Goldberg, , and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225.
- Levy, Omer and Yoav Goldberg. 2014. Neural word embeddings as implicit matrix factorization. In *Proceedings of NIPS*.
- Lewis, Mike and Mark Steedman. 2013. Combined distributional and logical semantics. *Transactions of the Association for Computational Linguistics*, 1:179–192.
- Maillard, Jean and Stephen Clark. 2015. Learning adjective meanings with a tensor-based skip-gram model. In *Proceedings of the SIGNLL Conference on Computational Natural Language Learning (CoNLL 2015)*, pages 327–331, Beijing, China.
- Maillard, Jean, Stephen Clark, and Edward Grefenstette. 2014. A type-driven tensor-based semantics for CCG. In *Proceedings of the EACL 2014 Type Theory and Natural Language Semantics Workshop (TTNLS)*, Gothenburg, Sweden.
- Manning, Christopher D., Prabhakar Raghavan, and Hinrich Schutze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- Marantz, A. 1984. *On the nature of grammatical relations*. MIT Press, Cambridge, MA.
- Marelli, Marco, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. SemEval-2014 task 1: Evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment.
- Mikolov, T., W. Yih, and G. Zweig. 2013. Linguistic regularities in continuous space word representations. In *Proceedings of HLT-NAACL*, pages 746–751.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013.*, pages 3111–3119, Lake Tahoe, Nevada, United States.
- Milajevs, Dmitrijs, Dimitri Kartsaklis, Mehrnoosh Sadrzadeh, and Matthew Purver. 2014. Evaluating neural word representations in tensor-based compositional settings. In *Proceedings of EMNLP*, Doha, Qatar.
- Miller, George A. 1995. WordNet: A lexical database for English. *Commun. ACM*, 38(11):39–41, November.
- Minnen, Guido, John Carroll, and Darren Pearce. 2001. Applied morphological processing of English. *Natural Language Engineering*, 7(3):207–223.
- Mitchell, J. and M. Lapata. 2008. Vector-based models of semantic composition. In *Proceedings of ACL*, pages 236–244, Columbus, OH.

- Mitchell, Jeff and Mirella Lapata. 2010. Composition in distributional models of semantics. *Cognitive Science*, 34:1388–1439.
- Montague, R. 1974. English as a formal language. In R. Montague, editor, *Formal Philosophy*. Yale University Press, New Haven, CT.
- Padó, Sebastian and Yves Peirsman. 2011. Introduction. In *Proceedings of the EMNLP 2011 Geometrical Models for Natural Language Semantics (GEMS 2011) Workshop*, page iii, Edinburgh. <https://sites.google.com/site/geometricalmodels/shared-evaluation>.
- Paperno, Denis, Nghia The Pham, and Marco Baroni. 2014. A practical and linguistically-motivated approach to compositional distributional semantics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 90–99.
- Pham, N., R. Bernardi, Y.-Z. Zhang, and M. Baroni. 2013. Sentence paraphrase detection: When determiners and word order make the difference. In *Proceedings of Towards a Formal Distributional Semantics Workshop at IWCS*, pages 21–29.
- Plate, Tony A. 1991. Holographic reduced representations: Convolution algebra for compositional distributed representations. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 30–35, Sydney, Australia.
- Plate, Tony A. 2000. Analogy retrieval and processing with distributed vector representations. *Expert Systems*, 17(1):29–40.
- Polajnar, Tamara and Stephen Clark. 2014. Improving distributional semantic vectors through context selection and normalisation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, Gothenburg, Sweden.
- Polajnar, Tamara, Luana Fagarasan, and Stephen Clark. 2014. Reducing dimensions of tensors in type-driven distributional semantics. In *Proceedings of EMNLP 2014*, Doha, Qatar.
- Polajnar, Tamara, Laura Rimell, and Stephen Clark. 2014. Evaluation of simple distributional compositional operations on longer texts. In *Proceedings of the 9th Language Resources and Evaluation Conference (LREC 2014)*, Reykjavik, Iceland.
- Polajnar, Tamara, Laura Rimell, and Stephen Clark. 2015. An exploration of discourse-based sentence spaces for compositional distributional semantics. In *Proceedings of the Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics (LSDSem)*, Lisbon, Portugal.
- Pollack, J. 1990. Recursive distributed representations. *Artificial Intelligence*, 46:77–105.
- Reddy, Siva, Diana McCarthy, and Suresh Manandhar. 2011. An empirical study on compositionality in compound nouns. In *Proceedings of The 5th International Joint Conference on Natural Language Processing 2011 (IJCNLP 2011)*, Chiang Mai, Thailand.
- Rimell, Laura, Stephen Clark, and Mark Steedman. 2009. Unbounded dependency recovery for parser evaluation. In *Proceedings of EMNLP*, pages 813–821, Singapore.
- Sadrzadeh, Mehrnoosh, Stephen Clark, and Bob Coecke. 2013. The Frobenius anatomy of word meanings I: subject and object relative pronouns. *Journal of Logic and Computation*, 23:1293–1317.
- Schütze, Hinrich. 1998. Automatic word sense discrimination. *Computational Linguistics*, 24(1):97–124.
- Smolensky, Paul. 1990. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46:159–216.
- Smolensky, Paul and Geraldine Legendre. 2006. *The Harmonic Mind: from neural computation to optimality-theoretic grammar*. MIT Press, Cambridge, MA.
- Socher, Richard, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1201–1211, Jeju, Korea.
- Socher, Richard, Christopher D. Manning, and Andrew Y. Ng. 2010. Learning continuous phrase representations and syntactic parsing with recursive neural networks. In *Proceedings of NIPS Deep Learning and Unsupervised Feature Learning Workshop*, Whistler, BC, Canada.
- Socher, Richard, Alex Perelygin, Jean Wu, Jason Chuang, Chris Manning, Andrew Ng, and Chris Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*, Seattle, USA.
- Steedman, Mark. 2000. *The Syntactic Process*. The MIT Press, Cambridge, MA.
- Turney, Peter D. and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37:141–188.

- Utsumi, Akira. 2012. Extending and evaluating a multiplicative model for semantic composition in a distributional semantic model. In *Proceedings of the 11th International Conference on Cognitive Modeling (ICCM-2012)*, pages 243–248.
- Vecchi, E., M. Baroni, and R. Zamparelli. 2011. (Linear) maps of the impossible: Capturing semantic anomalies in distributional space. In *Proceedings of the DISCO (Distributional Semantics and Compositionality) Workshop at ACL*, pages 1–9.
- Zanzotto, Fabio Massimo, Lorenzo Ferrone, and Marco Baroni. 2015. When the whole is not greater than the combination of its parts: A “decompositional” look at compositional distributional semantics. *Computational Linguistics*, 41(1).