

Representing Syntax and Composition with Geometric Transformations

Lorenzo Bertolini Julie Weeds David Weir Qiwei Peng

University of Sussex

Brighton, UK

{l.bertolini, j.e.weeds, d.j.weir, qiwei.peng}@sussex.ac.uk

Abstract

The exploitation of syntactic graphs (SyGs) as a word’s context has been shown to be beneficial for distributional semantic models (DSMs), both at the level of individual word representations and in deriving phrasal representations via composition. However, notwithstanding the potential performance benefit, the syntactically-aware DSMs proposed to date have huge numbers of parameters (compared to conventional DSMs) and suffer from data sparsity. Furthermore, the encoding of the SyG links (i.e., the syntactic relations) has been largely limited to linear maps. The knowledge graphs’ literature, on the other hand, has proposed light-weight models employing different geometric transformations (GTs) to encode edges in a knowledge graph (KG). Our work explores the possibility of adopting this family of models to encode SyGs. Furthermore, we investigate which GT better encodes syntactic relations, so that these representations can be used to enhance phrase-level composition via syntactic contextualisation.

1 Introduction

Representing words in terms of their syntactic co-occurrences has been long proposed, both for count-based (Padó and Lapata, 2007; Weir et al., 2016), and neural (Hermann and Blunsom, 2013; Levy and Goldberg, 2014; Komninos and Manandhar, 2016; Czarnowska et al., 2019; Vashishth et al., 2019) models of word meaning. Tested on benchmark word similarity tasks, such models often perform favourably to models based on proximal co-occurrence, particularly when the similarity or substitutability of two words is considered rather than their relatedness (Levy and Goldberg, 2014). However, the real promise of distributional models based on syntactic rather than proximal co-occurrence, is the potential for carrying out syntax-sensitive composition. For example, in the

Anchored Packed Tree (APT) model (Weir et al., 2016) lexemes, phrases, and sentences are represented as collections of typed occurrences, and composition is carried out by contextualising each element in its syntactic role. This leads to syntax-sensitive representations for phrases. For example, *glass window* and *window glass* have different representations due to the different syntactic roles played by each constituent.

Alongside count-based models, a variety of neural ones have been proposed to encode syntactic structure, focusing on different depths of the graph (Levy and Goldberg, 2014; Komninos and Manandhar, 2016; Marcheggiani and Titov, 2017; Vashishth et al., 2019; Emerson, 2020)). Of particular note here, Levy and Goldberg (2014) and Komninos and Manandhar (2016) each proposed models (DEP and EXT, respectively) which learn from local dependency relations, by extending the Skip-Gram with Negative sampling (SGNS) architecture from *word2vec* (Mikolov et al., 2013). Given a tuple of (*target, context*) words, e.g. (*rain, like*), a standard SGNS model can be trained to encode the probability of it being a true or a randomly sampled tuple. DEP and EXT, on the other hand, make use of both standard and syntactically contextualised tuples e.g., (*rain, \overline{dobj} , like*)¹. Whilst DEP was tested solely on word similarity tasks, Komninos and Manandhar (2016) applied large neural architectures to sentence level tasks and were thus able to demonstrate a positive impact of applying an additive composition strategy to syntax-aware representations.

There is of course an explosion in the number of parameters to be learnt in both DEP and EXT due to the many possible word-relation combinations which form the target vocabulary for these models (see Table 1). A possible solution, pro-

¹ \overline{dobj} indicating the inverse of the *dobj* relation

posed by [Czarnowska et al. \(2019\)](#), is the Dependency Matrix (DM) model which uses linear maps in the form of square matrices to encode relations. Here, the training objective is changed from predicting $(target, context)$ pairs to $(target, relation, context)$ triples, e.g., $(rain, \overline{dobj}, like)$. This model produced comparable results with DEP and EXT at the word level. Furthermore, compositional experiments on short phrases, specifically relative clauses, produced encouraging results when using the learned transformations. Yet, despite considerably reducing the number of parameters, this model still makes use of large word spaces and the square linear map is still costly to train.

Model	Learnable Parameters
DEP	223M
DM	51.6M
MuRE	21.5M
RotE	21.5M
RefE	21.5M
AttE	21.6M

Table 1: Learnable parameters for each model, given the same word (72k) and relation (88) vocabularies from the `text8` (parsed) corpus, and vector size of $n=300$.

The reformulation of the SGNS objective introduced by DM (i.e., moving from $(target, context)$ tuples to $(target, relation, context)$ triples) closely resembles a common practice in the knowledge graphs (KGs) literature (e.g. ([Trouillon et al., 2017](#); [Balazevic et al., 2019](#); [Chami et al., 2020](#))). Here, large, mainly factual, graphs are fed to neural models in the form of $(head, relation, tail)$. Compared to the syntactically-aware DSMs discussed above, many of the models proposed to encode KGs make use of a substantially lower number of parameters to encode both word and relations, as shown in Table 1. Furthermore, in order to represent the heterogeneous types of relations in KGs, researchers have experimented with models based on different types of geometric transformations (GTs). These include, but are not limited to, stretch ([Balazevic et al., 2019](#)), rotation ([Sun et al., 2019](#); [Chami et al., 2020](#)), reflection ([Chami et al., 2020](#)) and attention ([Chami et al., 2020](#)). However, in the KG literature, limited attention has been paid to the compositional nature of phrases. Single-token oriented vocabularies (where *New York* is represented by *New_York*), used in most KGs, work well for real-world entities, such as people or cities, but are prob-

lematic when considering compositional phrases such as *small cake*. As discussed by [Toutanova et al. \(2015\)](#), treating these phrases in the same way forces the vocabulary to grow immensely, and prevents the model from reasoning over new phrases in a compositional fashion. Hence, developing successful composition strategies is of interest to the KG community as well as more widely in Natural Language Inference (NLI).

Given the success that DM and other models have obtained in modelling syntax and syntactically driven composition, we propose to overcome the parameter and word-relation vocabulary problems by using GT models to encode syntactic graphs. We focus our investigation on four state of the art models from the knowledge-graphs literature, namely MuRE ([Balazevic et al., 2019](#)), and the three GTs-based models proposed by [Chami et al. \(2020\)](#): RotE, RefE and AttE. Despite the simplicity, MuRE has obtained competitive results, when compared to more complex models ([Chami et al., 2020](#))). Rotation has been used to model composition of relation representations ([Sun et al., 2019](#)). Attention has been frequently proposed as a plausible mechanism for composition (e.g. [Hudson and Manning \(2018\)](#); [Tay et al. \(2019\)](#); [Yin et al. \(2020\)](#); [Russin et al. \(2020\)](#)), whilst reflection is relatively under-studied ([Chami et al., 2020](#)). Furthermore, as discussed in Section 3, these models allow for an interesting comparison, as they can be grouped into three categories: tail modifiers (DM), head modifiers (RotE, RefE, AttE), and full modifiers (MuRE). Hence, we explore some of the transformational properties required to enable the successful encoding of syntactic relations, where success is defined in terms of their potential to support phrasal composition.

Our contributions are as follows. First, we show how lighter-weight models based on GTs can be used to encode both word and syntactic relations, frequently outperforming DM both in word similarity and compositional benchmarks. Second, for each model, we propose a tailored composition strategy, based on syntactic contextualisation of one (or more) of the phrase constituents. We hence show how to exploit the learned syntactic representations for composition, by comparing syntax-driven strategies for composition with simple addition. Third, we provide an analysis of which type of GTs better encode relations for syntactic contextualisation and enhanced composition.

2 Related Work

Knowledge graphs are complex data structures where nodes are concepts or entities (usually content words like *dog* or *Campari*) and edges are relations (e.g. *is_a*, *produced_in*) connecting entities to one another (e.g. *dog is_a mammal*, *Campari produced_in Italy*). Table 2 reports the number of distinct entities, relations and triples for three of the most investigated KGs, namely, FB15k-237 (Toutanova and Chen, 2015) YAGO3-10 (Mahdisoltani et al., 2015), and WN18RR (Dettmers et al., 2018), as well as a syntactic graph (SyG) constructed from the parsed corpus `text8`. The way these graphs are structured can vary significantly. Chami et al. (2020) showed how, among the presented KGs, only WN18RR has a significantly hierarchical structure.

Dataset	entities	relations	triples	graph type
WNRR18	31k	11	87k	KG
FB15k-237	15k	237	272k	KG
YAGO3-10	123k	33	1M	KG
<code>text8</code>	72k	88	12M*	SyG

Table 2: Statistics for the training splits of different datasets (* number of unique items, with observed repetitions, items raise to 18M).

Research on models for representing KGs has mainly focused on the ability to predict new connections between existing nodes. To overcome the problem of testing items that do not occur in the training set, many models have adopted negative sampling (NS) strategies in the training phase. The vocabulary of KG datasets is also largely single-token oriented. Models able to handle multi-token items have been proposed (Toutanova et al., 2015, 2016; Sun et al., 2019), but they focus on the composition of relations rather than entities, e.g., how a complex relation such as *married_to:son_of* might be split into multiple constituents and composed. Also relevant, Toutanova and Chen (2015) showed how syntax-augmented triples extracted from documents (e.g. (*Obama*, *nsubj:born_in:obj,USA*)) can be beneficial for KGs models, but did not investigate representing syntax or composition via embeddings.

Previous works (e.g. (Marcheggiani and Titov, 2017; Vashishth et al., 2019)) showed how SyGs could be encoded via graph convolutional networks (GCN) (Kipf and Welling, 2017). These large mod-

els are able to encode larger graphs (up to the sentence level), via sequences of convolutions along the edges of the graph. Such convolutions are frequently relation-specific and are also encoded via square matrices.

3 Theoretical Approach

In both the semantic (KG) and syntactic (SyG) domain, the starting point is typically a dataset D of positive triples (h, r, t) , with $h, t \in V = \{1, \dots, |V|\}$ and $r \in R = \{1, \dots, |R|\}$, where V and R are the sets of the indexes for the vocabulary of entities / words and relations, respectively. In both domains, the shared goals are: i) map entities $v \in V$ to embeddings e_v where $e \in \mathbb{R}^{|V| \times n}$, n being the dimensionality of the vectors; ii) map relations $r \in R$ in one – or more – space $\mathbb{R}^{|R| \times *}$. In this work, we focus on constructing a syntactic dataset of positive training triples from a corpus as in Czarnowska et al. (2019). All of the models we investigate rely on a negative sampling mechanism that generates a dataset D' of false triples. Each model was presented in its own original work with a tailored way to generate D' . Unless otherwise stated, we make use of the original mechanism.

As already discussed, we are interested in both word level and compositional level evaluation. Testing at the word level, e.g., using word similarity benchmarks, simply requires extraction of the word embeddings. Compositional tests, on the other hand, also require syntactic analysis of the phrase and extraction and application of the relation embeddings. The first step, is to generate a parsed version of the phrase. For example, syntactic analysis of the phrase *pour tea* will produce the root-as-head (Rh) (h, r, t) triple $(pour, \overline{obj}, tea)$, and the root-as-tail (Rt) (h, r, t) triple $(tea, \overline{obj}, pour)$. Such duplicity of representations was handled in DM by obtaining both representations and then summing the cosine similarities obtained when comparing each of the two representations with a given target. Whilst reasonably effective in the DM evaluation, this does not provide a single phrase-level representation and would become unwieldy for longer phrases and sentences. Weir et al. (2016) argued in favour of considering the syntactic root as the main element of any multi-token linguistic item. In our example, to compare *pour tea* with *drink water*, this would require us to consider the syntactic root in the context of its dependent i.e., how similar is the verb *pour* when contextualised by the direct

object *tea* to the verb *drink* when contextualised by the `direct_object` *water*? In models which modify the head of the triple (e.g., (Chami et al., 2020)), this would correspond to using the root-as-tail (Rt) analysis of the phrase. Here, we compare the two strategies empirically. Further, inspired by the growing success of (very large) bi-directional models such as ELMo (Peters et al., 2018) and BERT (Devlin et al., 2018) and also by recent evidence from the neuroscientific literature (Mollica et al., 2020; Fedorenko et al., 2020), suggesting that sentence processing strongly relies on identifying and composing smaller units of meaning, such as phrases, regardless of order of their constituents, we also propose a third compositional strategy which is bi-directional in nature. Here, the phrase-level representation is the sum of the root-as-head and the root-as-tail representations, making it more agnostic to the direction of the relation as well as the word order. However, phrases with different structures such as *glass window* and *window glass* will still have different representations due to the different roles played by each word in each relation.

In summary, we propose and investigate three different syntax-aware (*syn*) composition strategies: *syn-Rh* and *syn-Rt*, different solely in where the root is placed in the (*head, relation, tail*) triple; and *syn-BiD* (for bi-directional), constructed by adding the representations obtained by *syn-Rh* and *syn-Rt*. We now describe in detail the models investigated, together with our tailored *syn* composition strategy for each of them.

DM This model is an extension of SGNS, where a linear map, in the form of a $n \times n$ matrix, projects a word from the context space (e') into the target space (e), as in Equation 1:

$$u = e_h^T \cdot (W_r e_t') \quad (1)$$

where $e, e' \in \mathbb{R}^{|\mathcal{V}| \times n}$, and $W \in \mathbb{R}^{|\mathcal{R}| \times n \times n}$. Since the tail word is projected into the space occupied by the head word, we refer to this model as a tail-modifier. u is then used to compute standard SGNS loss (Equation 2):

$$\sum_{(h,r,t) \in D} \log \sigma(u) + \sum_{(h,r,t) \in D'} \log \sigma(-u) \quad (2)$$

Phrase representations will be constructed following our three syntactic composition strategies. As

a baseline, common to all models, we use addition (*add*) of the queried head and tail entities embeddings, as in Equation 3²:

$$e_{add} = e_h + e_t \quad (3)$$

We propose *syn* composition for the DM model to be obtained via u (Equation 1), as in Equation 4:

$$e_{syn} = e_h + (W_r e_t') \quad (4)$$

MuRE This architecture falls into the family of translation models (Chami et al., 2020). Here, both the entities go through a transformation and so we refer to this model as a full-modifier. The tail entity is shifted with a translation (i.e. offset), and a stretch, in the form of a $n \times n$ diagonal matrix, is applied to the head entity. Embeddings are then fed to a distance function $d(x, y) = \|x - y\|$ and the model minimises the Bernoulli negative log-likelihood loss, using Equation 5, to estimate the probability of the triple being from D :

$$p(h, r, l) = \sigma(-d(W_r e_h, e_t + w_r)^2 + b_h + b_t) \quad (5)$$

Here, $W \in \mathbb{R}^{|\mathcal{R}| \times n \times n}$ contains $|\mathcal{R}|$ diagonal matrices (each corresponding to a relation-specific stretch), $w \in \mathbb{R}^{|\mathcal{R}| \times n}$ hosts $|\mathcal{R}|$ translation vectors, and $b \in \mathbb{R}^{|\mathcal{V}| \times n}$ the entity biases. Again, additive composition is carried out by adding the queried embedding for the phrase’s constituents. Syntactic composition is implemented by adapting the model’s score function (Equation 6):

$$e_{syn} = W_r e_h + (e_t + w_r) \quad (6)$$

RotE, RefE These models optimise a full cross-entropy loss. Like MuRE, square distance between two vectors is used as a score function. Unlike the previous model, they apply a Givens rotation (Rot) or reflection (Ref), as defined in Chami et al. (2020), and a translation to the head entity. Thus, we refer to these models as head-modifiers. Syntactic composition is defined via the score functions in Equations 7 and 8:

$$e_{syn} = (\text{Rot}(T_r) e_h + t_r) + e_t \quad (7)$$

$$e_{syn} = (\text{Ref}(F_r) e_h + f_r) + e_t \quad (8)$$

where $T, F \in \mathbb{R}^{|\mathcal{R}| \times \frac{n}{2}}$ each contain $|\mathcal{R}|$ diagonal matrices (each corresponding to a relation-specific Givens rotation or reflection), and $t, f \in \mathbb{R}^{|\mathcal{R}| \times n}$ are relation-specific translations.

²This corresponds to simple-sum composition in the original work by Czarnowska et al. (2019).

AttE Intuitively, AttE is designed to model the contribution of different GTs (in this case just rotation and reflection). This is achieved via a self-attention mechanism. Given two embeddings x , y , and an attention vector a , attention scores are computed via Equation 9:

$$(\alpha_x, \alpha_y) = \text{Softmax}(a^T x, a^T y) \quad (9)$$

These scores are then averaged (Equation 10):

$$\text{Att}(x, y; a) = (\alpha_x x + \alpha_y y) \quad (10)$$

To actively select the most suitable transformation for a given triple, rotation and reflection are applied to the head-entity embedding (Equation 11):

$$q_{\text{Rot}} = \text{Rot}(T_r)e_h, \quad q_{\text{Ref}} = \text{Ref}(F_r)e_h \quad (11)$$

The two representations are then combined using a self attention mechanism (Equation 12):

$$Q(h, r) = \text{Att}(q_{\text{Rot}}, q_{\text{Ref}}; a_r) + p_r \quad (12)$$

with $p \in \mathbb{R}^{|R| \times n}$ as the relation-specific translation. Q and the e_t are then used as arguments for d as in Equation 5. Syntactically contextualised composition (*syn*) for AttE is implemented via Equation 13:

$$e_{\text{syn}} = Q(h, r) + e_t \quad (13)$$

4 Experiments

Our main aim is to investigate the potential of models in terms of constructing high quality word representations and their support for composition. To this end, experiments were carried out with a set of models trained on KGs, and a second set of models trained on SyGs. This allows us to investigate the value of encoding distributional information from SyGs or whether KGs alone might be a sufficient source of data to obtain competitive results. We hypothesise that when using KGs alone: i) word similarity tasks might yield high results; ii) compositional evaluation will yield poor results. As for models trained on SyG, we expect to see: i) a generally improved performance on most tasks, when compared to models trained on KGs; ii) larger models to be penalised across benchmarks and for syntactically-contextualised (*syn*) composition.

4.1 Experimental setup

Benchmarks We divide our quantitative experiments between word similarity and composition

tasks. For the word similarity tasks, we focus on SimLex (Hill et al., 2015), MEN (Bruni et al., 2014), and both similarity (WS_s) and relatedness (WS_r) split of the WordSim353 (Finkelstein et al., 2001) datasets. For every word pair, we produce a model’s prediction using cosine similarity (CS). We compare model predictions and human judgements using Spearman’s ρ .

For the compositional investigation, we focus on the Mitchell and Lapata (2010) (ML10) dataset. Items in this benchmark consist of pairs of two-token phrases (e.g. (*pour tea-drink water*)) paired with human judgements on their similarity. Phrases are composed using the four different presented strategies and the obtained representations are compared via CS. Again, CS and human ratings are compared via ρ . We selected this benchmark for two main reasons: i) the models’ structures lend themselves straightforwardly to syntactically contextualised (*syn*) composition strategies for a two-token item³; ii) the dataset is pre-split into three syntactic-relation classes (i.e. adjective-nouns (AN), verb-objects (VO) and noun-nouns (NN)) and this division offers an opportunity for a more in-depth investigation on how different models and operations manage to embed different syntactic relations.

We trained each set of models with three random initialisation, and report the mean and standard error (SE) of the obtained ρ s.

Implementation For MurE, RotE, RefE and AttE we adapt the original PyTorch code. Since an official release of the DM is not available, we implemented a PyTorch version of the model⁴.

We trained the first set of GT models on the WN18RR dataset, tuning negative sampling rate (NS), optimiser and learning rate using mean reciprocal rank (MRR) on the development set⁵. Epochs were kept stable at 50 and n to 300. We focused on WN18RR as YAGO3-10 shares a minimal vocabulary with the selected word-similarity and compositional benchmarks. FB15k-237, on the other hand, has all the entities encrypted. The models obtained from this training set were then evaluated on both word-similarity and compositional tasks (see Table 3) to provide a baseline for the SyG models.

³Czarnowska et al. (2019) proposed a more complex composition strategy, specifically for relative clause phrases which we do not consider here.

⁴https://github.com/lorenzoscottb/findings_ACL2021

⁵using the dataset’s original splits.

	Simlex	MEN	WS_s	WS_r	Adjective Nouns	Verb Objects	Noun-Noun
MuRE	.38±.01	.45±.00	.42±.01	.21±.03	.19±.03	.31±.00	.13±.01
RotE	.35±.01	.54±.00	.59±.00	.30±.02	.18±.03	.33±.00	.20±.02
RefE	.36±.01	.54±.00	.57±.00	.30±.01	.16±.04	.37±.01	.14±.02
AttE	.36±.01	.54±.00	.58±.01	.29±.00	.20±.00	.32±.00	.18±.00

Table 3: Spearman ρ s’ (mean \pm SE) obtained on all selected benchmarks, for knowledge-graph models trained on WN18RR dataset.

A second set of models was trained on the `text8`⁶ corpus, parsed with spaCy (Honnibal and Johnson, 2015). Following Czarnowska et al. (2019), minimum item count, epochs, NS, optimiser and learning rate were fine-tuned on SimLex. Hyperparameters are selected from the union of the ones proposed in (Balazevic et al., 2019; Czarnowska et al., 2019; Chami et al., 2020). All the models share the same number of dimensions, i.e., $n = 300$. For a fair comparison, all experiments for this set have been conducted on the vocabulary shared across the models. Final coverage and best hyperparameters are reported in Appendix A.2 and A.1. All models were trained using NVIDIA Titan V GPUs.

4.2 Results

WN18RR trained models We begin our quantitative investigation evaluating models from the knowledge graph literature, trained on WN18RR, on all benchmarks. Looking at Table 3, we note that these models, compared to models trained on `text8` or similar distributional models trained on much larger corpora, achieve competitive results on the word similarity benchmarks, especially in the historically challenging SimLex dataset, despite the small vocabulary and training samples.

A possible explanation for these results lies in how entities co-occur in the training data. First of all, WN18RR has a limited vocabulary (see Table 2), and is poorly populated by adjectives. Furthermore, noun and verbs, two part of speech (POS) that frequently co-occur between each other in natural language, here mainly occur within each other (i.e. verb with verb, noun with noun). In few cases, especially for verbs, the co-occurrences are not only limited to the same POS, but interest the very same word. All models perform much worse on the relatedness split of WS-353 than the similarity split. This might be expected, for models trained on WordNet data. As predicted, the performance is generally poor for composition benchmarks. An

⁶<http://matmahoney.net/dc/textdata>

exception seems to be the VO subset, where models achieve results that, as will be presented shortly, are competitive also for `text8`-trained models.

Word similarity Our motivation for experiments with models trained on `text8` is to understand whether models previously proposed for representing KGs are competitive with distributional models such as DM in their ability to embed word and syntactic relations. Results for word-similarity are presented in Table 4.

	Simlex	MEN	WS_s	WS_r
DM	.12±.01	.60±.01	.59±.02	.51±.03
MuRE	.17±.01	.64±.00	.69±.01	.58±.00
RotE	.17±.00	.64±.00	.70±.00	.58±.01
RefE	.18±.01	.63±.01	.70±.01	.56±.00
AttE	.16±.00	.61±.00	.69±.01	.57±.01

Table 4: Spearman ρ s’ (mean \pm SE) obtained on word-similarity benchmarks, with models trained on `text8` corpus.

First, scores on SimLex are much lower than: i) those achieved by the KG-trained models; ii) those presented elsewhere for DM in the literature (Czarnowska et al., 2019). We note that the corpus we used to train the models is significantly smaller than the one used to train DM by the original authors, and we assume that this, combined with the low frequency of SimLex items in our corpus, is the main reason for these differences. Results for DM on the other word similarity benchmarks are much closer to the performance achieved by the original authors and, on these benchmarks, DM clearly outperforms the baseline of models trained on WN18RR. However, most notably, GT models trained on the same data as DM, not only achieve comparable results to DM, but they almost always outperform it, both in similarity-based and relatedness-based benchmarks. Moreover, DM seems to show the highest variation, especially for WN_s and WN_r.

Composition Table 5 shows the results for all `text8`-trained models on the compositional

		Adjective Nouns	Verb Objects	Noun-Nouns	Average
DM	<i>add</i>	.39±.02	.31±.03	.43±.03	.37±.02
	<i>syn-Rh</i>	.26±.02	.18±.02	.25±.03	.23±.02
	<i>syn-Rt</i>	.32±.03	.14±.02	.20±.02	.22±.02
	<i>syn-BiD</i>	.33±.02	.14±.02	.34±.03	.27±.02
MuRE	<i>add</i>	.47±.01	.35±.01	.40±.00	.41±.00
	<i>syn-Rh</i>	.51±.01	.34±.01	.44±.01	.43±.00
	<i>syn-Rt</i>	.49±.01	.36±.01	.43±.01	.43±.01
	<i>syn-BiD</i>	.49±.01	.36±.01	.46±.01	.44±.00
RotE	<i>add</i>	.49±.00	.37±.00	.43±.00	.43±.00
	<i>syn-Rh</i>	.48±.01	.36±.01	.41±.01	.42±.01
	<i>syn-Rt</i>	.47±.02	.35±.01	.41±.01	.41±.00
	<i>syn-BiD</i>	.49±.00	.38±.00	.45±.01	.44±.00
RefE	<i>add</i>	.48±.01	.36±.00	.43±.01	.42±.00
	<i>syn-Rh</i>	.49±.01	.36±.01	.43±.01	.43±.01
	<i>syn-Rt</i>	.48±.01	.34±.02	.43±.01	.42±.01
	<i>syn-BiD</i>	.48±.00	.38±.01	.46±.01	.44±.01
AttE	<i>add</i>	.46±.01	.35±.01	.41±.01	.41±.00
	<i>syn-Rh</i>	.47±.01	.35±.00	.43±.01	.41±.01
	<i>syn-Rt</i>	.45±.01	.29±.01	.43±.01	.39±.00
	<i>syn-BiD</i>	.48±.02	.36±.00	.46±.00	.43±.00

Table 5: Spearman ρ_s (mean \pm SE) obtained on Mitchell and Lapata (2010) benchmark, with models trained on `text8` corpus. Phrasal composition is carried out by element-wise addition (*add*), and the three proposed syntax (*syn*) aware strategies: root as head (*syn-Rh*), root as tail (*syn-Rt*) and bidirectional (*syn-BiD*). **Best results for each Phrase Type.**

benchmark. Again, GT models show competitive results, and generally outperform DM, which fails at improving its performance with *syn* composition. This last evidence is reversed in all other models. That is, they all achieve best performance with one of the syntax-aware composition methods. Looking closer, we can see that, in most cases, the best *syn* method is the bi-directional one, with the exceptions of MUuRE, RotE and RefE’s AN phrases. Notably, *syn-BiD* is almost never a mere average of the two representations that originated it. In many cases, and especially for AttE, *syn-BiD* representations produce a significantly larger gain in performance, when compared to both *syn-Rt* and *syn-Rh*. From the single model perspective, the best performing one is RefE. Syntax-aware methods based on reflection always outperform the additive baseline, and also obtained the best score in the average sections, via bi-directional composition. Again, DM is the model showing the highest variation in results. This provides further evidences in favour of the lightweight models taken from the KG literature

4.3 Statistical Analysis

All correlations were tested for significance, adopting the Holm correction (Holm, 1979) to account for the large number of tests, and we observed no $p < .05$. As the main interest of our work was the

compositional investigation (reported in Table 5), a global comparison was conducted to test whether observed differences in correlations were also significant. We adopted a paired two-tail bootstrap analysis (Berg-Kirkpatrick et al., 2012; Søggaard et al., 2014; Dror et al., 2018), performed independently between results from the three seeds. Given the large number of comparisons, a Holm correction was adopted within the same Phrase Type. Results (see A.3 for more details) showed that, among all models, the only one that generated a number of insignificant differences was DM, mainly pertaining to different strategies for composing NN items.

4.4 Qualitative Analysis

We now investigate the impact of relation representations on word vectors and composition from a qualitative point of view. Here, we focus on the model that quantitative tests indicated as the most promising one: RefE. We will start at the word level, looking at syntactically contextualised single words. The interest here, is to see if clear relation-driven clusters can be identified within a reduced space. To do so, we contextualise the set of roots from ML10 (e.g. *amount* in *vast amount*), and reduce the dimensions through PCA. Results in Figure 2 suggest that the three syntactic relations adopted for contextualisation (i.e. *amod*, *dobj*,

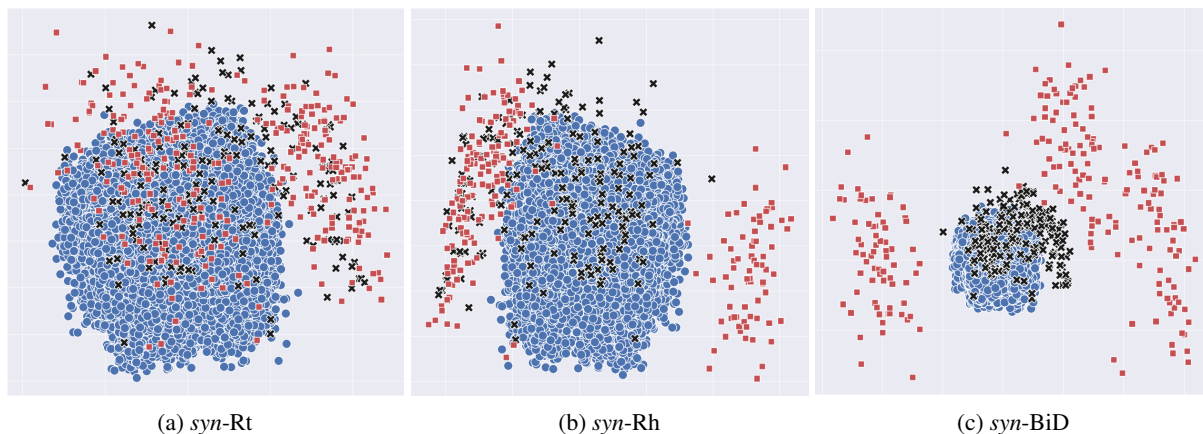


Figure 1: PCA visualisation of RefE vector space. Images show the same word (●) and *add*-composed vectors (×), in the context of representations composed with the four different syntax-aware (■) composition methods. All composed vectors represent the set of phrases from the Mitchell and Lapata (2010) benchmark.

nmod) appear to generate as many distinguishable clusters. Despite being limited, these results support evidence for syntactic subspace probed out of mBert (Chi et al., 2020).




Figure 2: PCA visualisation of syntactically contextualised root-items from ML10 phrases using RefE and reflection. AN roots are contextualised using *amod* (×), VO via *dobj* (●), NN via *nmod* (■).

Concluding, we explore how composition strategies behave with respect to the word representations. To do so, we concatenate representations obtained by *add*-composing the set of ML10 items with the full original space, and each syntax-aware strategy separately. The three obtained sets of concatenation (i.e. word–*add*–*syn*-Rt; word–*add*–*syn*-Rh; word–*add*–*syn*-BiD) is then independently reduced to $n=2$ through principal component analysis (PCA). Results are reported in Figure 1. As it can be observed throughout the three reductions,

and mostly in Figure 1c, phrase representations obtained via simple addition mainly lie within the perimeter of the word space. A similar pattern is observed in Figure 1a, with *syn*-Rt. Phrases composed by using the root as the head of the triple are still fairly close to the word-space perimeter, but tend to abandon its centre. Lastly, Figure 1c shows how bi-directional representations lie scattered fairly distant from the word and *add*-composed representations. This last observation is contrary to theories suggesting that representations at every level (word, phrase, sentence, etc..) should lie within the same space (e.g. Weir et al. (2016)). However, it may support recent work from neuroscience (e.g. Ding et al. (2016)) suggesting that the brain networks processing word, phrases and sentences do not completely overlap.

5 Discussion

Our results strongly suggest that light-weight models presented in the knowledge-graphs literature can be efficiently applied to syntactic-graphs, and be converted to distributional models that are consistently able to make use of the learned word and relation representations to improve semantic phrase-composition. From the model-theoretical point of view, evidence suggests that constraining linear maps with a reflection (together with a non-linear translation) seems to be the most efficient way of encoding syntactic relations. Our quantitative results also contribute to the debates on how sequential language data, or English at least, should be processed and what the role of syntactic information should be. As mentioned in

Section 3, the models selected distinguish between being tail (DM), head (RotE, RefE and AttE) and full (MuRE) modifiers. Further, we can change the syntactic focus of any of these models by adopting the *syn-Rt* composition strategy instead of the *syn-Rh* strategy. However, in our experiments, the head-modifier models (RotE, RefE and AttE) outperformed the tail-modifier and full models (DM and MuRE) and achieved a better results with the *syn-Rh* strategy than the *syn-Rt* strategy, i.e., when the syntactic root of the phrase was taken as the head of the triple rather than as the tail. In other words, it appears better to contextualise the root and compose with its dependent, which opposes the linguistic arguments put forward by Weir et al. (2016). However, even more notably, the *syn-BiD* composition strategy, which combines the *syn-Rh* and *syn-Rt* representations, generally gave a further boost to performance. This is further evidence that bi-directional information is more informative than uni-directional information, not just in large neural models such as LSTMs and transformers, and supports recent theory from neuroscience which argues that what is crucial for composition is not the overall structure nor the root, but that we can identify a phrase’s constituents and the relation they have (Mollica et al., 2020). Evidence in favour of the fact that composition strongly relies on local dependencies based on syntactic structure was also found by Saphra and Lopez (2020). Such work suggests that LSTMs learn to compose following a hierarchical structure, driven by syntax, and that they rely on the learned short sequences to build longer and more reliable ones. Taken altogether, the evidence from different language-related fields is becoming more compelling that syntax and phrase composition should play an important role in the composition of larger units of meaning.

6 Conclusions and Further Work

We have shown how GT models previously proposed for encoding KGs can be adapted to encode syntactic information in a distributional model. We have demonstrated the high quality nature of the distributional word representations and the potential for using syntactically-contextualised composition strategies for phrases. In particular, we have demonstrated the competitiveness of lighter-weight GT models when compared to more general models based solely on unconstrained linear maps, such as DM. Further, our analysis has shown how learned

representations for syntactic relations can be efficiently exploited at the word level, transforming a word through part-of-speech related regions of the space, and at the phrase level, generating superior composed representations. Furthermore, we have shown, among the different GTs, reflection seems to be the most promising for encoding syntactic relations. Future work will focus on composition on larger scale, syntactic-relation composition, and whether syntactic and semantic graph can be simultaneously embedded using this framework.

Acknowledgements

This research was supported by EPSRC grant no. 2129720: *Composition and Entailment in Distributed Word Representations*. We also thank the anonymous reviewers for their helpful comments, and NVIDIA for the donation of the GPU that supported our work. The first author would also like to thank Gabriele Paveri, for the years of conversations on human language, and constantly doubting the author’s ideas.

References

- Ivana Balazevic, Carl Allen, and Timothy Hospedales. 2019. [Multi-relational poincaré graph embeddings](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 4463–4473. Curran Associates, Inc.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. [An empirical investigation of statistical significance in NLP](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 995–1005, Jeju Island, Korea. Association for Computational Linguistics.
- Elia Bruni, Nam Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.
- Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. 2020. [Low-dimensional hyperbolic knowledge graph embeddings](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6901–6914, Online. Association for Computational Linguistics.
- Ethan A. Chi, John Hewitt, and Christopher D. Manning. 2020. [Finding universal grammatical relations in multilingual BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5577, Online. Association for Computational Linguistics.

- Paula Czarnecka, Guy Emerson, and Ann Copestake. 2019. Words are vectors, dependencies are matrices: Learning word embeddings from dependency graphs. In *Proceedings of the 13th International Conference on Computational Semantics - Long Papers*, pages 91–102, Gothenburg, Sweden. Association for Computational Linguistics.
- Tim Dettmers, Minervini Pasquale, Stenetorp Pontus, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the 32th AAAI Conference on Artificial Intelligence*, pages 1811–1818.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ArXiv e-prints*.
- N. Ding, L. Melloni, H. Zhang, Xing Tian, and D. Poeppel. 2016. Cortical tracking of hierarchical linguistic structures in connected speech. *Nature neuroscience*, 19:158 – 164.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. The hitchhiker’s guide to testing statistical significance in natural language processing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392, Melbourne, Australia. Association for Computational Linguistics.
- Guy Emerson. 2020. Autoencoding pixies: Amortised variational inference with graph convolutions for functional distributional semantics. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3982–3995, Online. Association for Computational Linguistics.
- Evelina Fedorenko, Idan Asher Blank, Matthew Siegelman, and Zachary Mineroff. 2020. Lack of selectivity for syntax relative to word meanings throughout the language network. *Cognition*, 203:104348.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: The concept revisited. In *Proceedings of the 10th International Conference on World Wide Web*, pages 406–414, New York, NY, USA. ACM.
- Karl Moritz Hermann and Phil Blunsom. 2013. The role of syntax in vector space models of compositional semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 894–904, Sofia, Bulgaria. Association for Computational Linguistics.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- S. Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6:65–70.
- Matthew Honnibal and Mark Johnson. 2015. An improved non-monotonic transition system for dependency parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1373–1378, Lisbon, Portugal. Association for Computational Linguistics.
- Drew Arad Hudson and Christopher D. Manning. 2018. Compositional attention networks for machine reasoning. In *International Conference on Learning Representations*.
- Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the 5th International Conference on Learning Representations, ICLR ’17*.
- Alexandros Komninos and Suresh Manandhar. 2016. Dependency based embeddings for sentence classification tasks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1490–1500, San Diego, California. Association for Computational Linguistics.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 302–308, Baltimore, Maryland. Association for Computational Linguistics.
- F. Mahdisoltani, J. Biega, and Fabian M. Suchanek. 2015. Yago3: A knowledge base from multilingual wikipedias. In *CIDR*.
- Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1506–1515, Copenhagen, Denmark. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Jeff Mitchell and Mirella Lapata. 2010. Composition in distributional models of semantics. *Cognitive Science*, 34(8):1388–1429.
- Francis Mollica, Matthew Siegelman, Evgeniia Diachek, Steven T. Piantadosi, Zachary Mineroff, Richard Futrell, Hope Kean, Peng Qian, and Evelina Fedorenko. 2020. Composition is the core driver of the language-selective network. *Neurobiology of Language*, 1(1):104–134.
- Sebastian Padó and Mirella Lapata. 2007. Dependency-based construction of semantic space models. *Computational Linguistics*, 33(2):161–199.

- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Jacob Russin, Jason Jo, Randall O’Reilly, and Yoshua Bengio. 2020. [Compositional generalization by factorizing alignment and translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 313–327, Online. Association for Computational Linguistics.
- Naomi Saphra and Adam Lopez. 2020. [LSTMs compose—and Learn—Bottom-up](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2797–2809, Online. Association for Computational Linguistics.
- Anders Søgaard, Anders Johannsen, Barbara Plank, Dirk Hovy, and Hector Martínez Alonso. 2014. [What’s in a p-value in NLP?](#) In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 1–10, Ann Arbor, Michigan. Association for Computational Linguistics.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. [RotatE: Knowledge graph embedding by relational rotation in complex space](#). In *International Conference on Learning Representations*.
- Yi Tay, Anh Tuan Luu, Aston Zhang, Shuohang Wang, and Siu Cheung Hui. 2019. [Compositional deattention networks](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 6135–6145. Curran Associates, Inc.
- Kristina Toutanova and Danqi Chen. 2015. [Observed versus latent features for knowledge base and text inference](#). In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*, pages 57–66, Beijing, China. Association for Computational Linguistics.
- Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. [Representing text for joint embedding of text and knowledge bases](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509, Lisbon, Portugal. Association for Computational Linguistics.
- Kristina Toutanova, Victoria Lin, Wen-tau Yih, Hoi-fung Poon, and Chris Quirk. 2016. [Compositional learning of embeddings for relation paths in knowledge base and text](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1434–1444, Berlin, Germany. Association for Computational Linguistics.
- Théo Trouillon, Christopher R. Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. 2017. [Knowledge graph completion via complex tensor factorization](#). *J. Mach. Learn. Res.*, 18:130:1–130:38.
- Shikhar Vashishth, Manik Bhandari, Prateek Yadav, Piyush Rai, Chiranjib Bhattacharyya, and Partha Talukdar. 2019. [Incorporating syntactic and semantic information in word embeddings using graph convolutional networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3308–3318, Florence, Italy. Association for Computational Linguistics.
- David Weir, Julie Weeds, Jeremy Reffin, and Thomas Kober. 2016. [Aligning packed dependency trees: a theory of composition for distributional semantics](#). *Computational Linguistics, special issue on Formal Distributional Semantics*, 42(4):727–761.
- Da Yin, Tao Meng, and Kai-Wei Chang. 2020. [SentiBERT: A transferable transformer-based architecture for compositional sentiment semantics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3695–3706, Online. Association for Computational Linguistics.
- Alexey Zobnin and Evgenia Elisratova. 2019. [Learning word embeddings without context vectors](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 244–249, Florence, Italy. Association for Computational Linguistics.