



HAL
open science

Composition of Embeddings : Lessons from Statistical Relational Learning

Damien Sileo, Tim van de Cruys, Camille Pradel, Philippe Muller

► **To cite this version:**

Damien Sileo, Tim van de Cruys, Camille Pradel, Philippe Muller. Composition of Embeddings : Lessons from Statistical Relational Learning. 8th Joint Conference on Lexical and Computational Semantics (SEM 2019), Jun 2019, Minneapolis, United States. pp.33-43. hal-02397476

HAL Id: hal-02397476

<https://hal.science/hal-02397476>

Submitted on 6 Dec 2019

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

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



Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in:
<http://oatao.univ-toulouse.fr/24995>

Official URL

DOI : <https://doi.org/10.18653/v1/s19-1004>

To cite this version: Sileo, Damien and Van de Cruys, Tim and Pradel, Camille and Muller, Philippe *Composition of Embeddings : Lessons from Statistical Relational Learning*. (2019) In: 8th Joint Conference on Lexical and Computational Semantics (SEM 2019), 6 June 2019 - 7 June 2019 (Minneapolis, United States).

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Composition of Sentence Embeddings: Lessons from Statistical Relational Learning

Damien Sileo^{3,1}, Tim Van De Cruys², Camille Pradel¹, and Philippe Muller³

¹Synapse Développement, Toulouse, France

²IRIT, CNRS, France

³IRIT, University of Toulouse, France

damien.sileo@synapse-fr.com

Abstract

Various NLP problems – such as the prediction of sentence similarity, entailment, and discourse relations – are all instances of the same general task: the modeling of semantic relations between a pair of textual elements. A popular model for such problems is to embed sentences into fixed size vectors, and use composition functions (e.g. concatenation or sum) of those vectors as features for the prediction. At the same time, composition of embeddings has been a main focus within the field of Statistical Relational Learning (SRL) whose goal is to predict relations between entities (typically from knowledge base triples). In this article, we show that previous work on relation prediction between texts implicitly uses compositions from baseline SRL models. We show that such compositions are not expressive enough for several tasks (e.g. natural language inference). We build on recent SRL models to address textual relational problems, showing that they are more expressive, and can alleviate issues from simpler compositions. The resulting models significantly improve the state of the art in both transferable sentence representation learning and relation prediction.

1 Introduction

Predicting relations between textual units is a widespread task, essential for discourse analysis, dialog systems, information retrieval, or paraphrase detection. Since relation prediction often requires a form of understanding, it can also be used as a proxy to learn transferable sentence representations. Several tasks that are useful to build sentence representations are derived directly from text structure, without human annotation: sentence order prediction (Logeswaran et al., 2016; Jernite et al., 2017), the prediction of previous and subsequent sentences (Kiros et al., 2015; Jernite et al., 2017), or the prediction of explicit dis-

course markers between sentence pairs (Nie et al., 2017; Jernite et al., 2017). Human labeled relations between sentences can also be used for that purpose, e.g. inferential relations (Conneau et al., 2017). While most work on sentence similarity estimation, entailment detection, answer selection, or discourse relation prediction seemingly uses task-specific models, they all involve predicting whether a relation R holds between two sentences s_1 and s_2 . This genericity has been noticed in the literature before (Baudiš et al., 2016) and it has been leveraged for the evaluation of sentence embeddings within the SentEval framework (Conneau et al., 2017).

A straightforward way to predict the probability of (s_1, R, s_2) being true is to represent s_1 and s_2 with d -dimensional embeddings h_1 and h_2 , and to compute sentence pair features $f(h_1, h_2)$, where f is a composition function (e.g. concatenation, product, ...). A softmax classifier g_θ can learn to predict R with those features. $g_\theta \circ f$ can be seen as a reasoning based on the content of h_1 and h_2 (Socher et al., 2013).

Our contributions are as follows:

- we review composition functions used in textual relational learning and show that they lack expressiveness (section 2);
- we draw analogies with existing SRL models (section 3) and design new compositions inspired from SRL (section 4);
- we perform extensive experiments to test composition functions and show that some of them can improve the learning of representations and their downstream uses (section 6).

2 Composition functions for relation prediction

We review here popular composition functions used for relation prediction based on sentence em-

beddings. Ideally, they should simultaneously fulfill the following minimal requirements:

- make use of interactions between representations of sentences to relate;
- allow for the learning of asymmetric relations (e.g. entailment, order);
- be usable with high dimensionalities (parameters θ and f should fit in GPU memory).

Additionally, if the main goal is transferable sentence representation learning, compositions should also incentivize gradually changing sentences to lie on a linear manifold, since transfer usually uses linear models. Another goal can be learning of transferable *relation* representation. Concretely, a sentence encoder and f can be trained on a base task, and $f(h_1, h_2)$ can be used as features for transfer in another task. In that case, the geometry of the sentence embedding space is less relevant, as long as the $f(h_1, h_2)$ space works well for transfer learning. Our evaluation will cover both cases.

A straightforward instantiation of f is concatenation (Hooda & Kosseim, 2017):

$$f_{[\cdot]}(h_1, h_2) = [h_1, h_2] \quad (1)$$

However, *interactions* between s_1 and s_2 cannot be modeled with $f_{[\cdot]}$ followed by a softmax regression. Indeed, $f_{[\cdot]}(h_1, h_2)\theta$ can be rewritten as a sum of independent contributions from h_1 and h_2 , namely $\theta_{[0:d]}h_1 + \theta_{[d:2d]}h_2$. Using a multi-layer perceptron before the softmax would solve this issue, but it also harms sentence representation learning (Conneau et al., 2017; Logeswaran et al., 2018), possibly because the perceptron allows for accurate predictions even if the sentence embeddings lie in a convoluted space. To promote interactions between h_1 and h_2 , element-wise product has been used in Baudiš et al. (2016):

$$f_{\odot}(h_1, h_2) = h_1 \odot h_2 \quad (2)$$

Absolute difference is another solution for sentence similarity (Mueller & Thyagarajan, 2016), and its element-wise variation may equally be used to compute informative features:

$$f_{-}(h_1, h_2) = |h_1 - h_2| \quad (3)$$

The latter two were combined into a popular instantiation, sometimes referred as *heuristic matching* (Tai et al., 2015; Kiros et al., 2015; Mou et al.,

2016):

$$f_{\odot-}(h_1, h_2) = [h_1 \odot h_2, |h_2 - h_1|] \quad (4)$$

Although effective for certain similarity tasks, $f_{\odot-}$ is symmetrical, and should be a poor choice for tasks like entailment prediction or prediction of discourse relations. For instance, if R_e denotes entailment and $(s_1, s_2) = (\text{“It just rained”}, \text{“The ground is wet”})$, (s_1, R_e, s_2) should hold but not (s_2, R_e, s_1) . The $f_{\odot-}$ composition function is nonetheless used to train/evaluate models on entailment (Conneau et al., 2017) or discourse relation prediction (Nie et al., 2017).

Sometimes $[h_1, h_2]$ is concatenated to $f_{\odot-}(h_1, h_2)$ (Ampomah et al., 2016; Conneau et al., 2017). While the resulting composition is asymmetrical, the asymmetrical component involves no interaction as noted previously. We note that this composition is very commonly used. On the SNLI benchmark,¹ 12 out of the 25 listed sentence embedding based models use it, and 7 use a weaker form (e.g. omitting f_{\odot}).

The outer product \otimes has been used instead for asymmetric multiplicative interaction (Jernite et al., 2017):

$$f_{\otimes}(h_1, h_2) = h_1 \otimes h_2 \text{ where } (h_1 \otimes h_2)_{i,j} = h_{1i}h_{2j} \quad (5)$$

This formulation is expressive but it forces g_{θ} to have d^2 parameters per relation, which is prohibitive when there are many relations and d is high.

The problems outlined above are well known in SRL. Thus, existing compositions (except f_{\otimes}) can only model relations superficially for tasks currently used to train state of the art sentence encoders, like NLI or discourse connectives prediction.

3 Statistical Relational Learning models

In this section we introduce the context of statistical relational learning (SRL) and relevant models. Recently, SRL has focused on efficient and expressive relation prediction based on embeddings. A core goal of SRL (Getoor & Taskar, 2007) is to induce whether a relation R holds between two arbitrary entities e_1, e_2 . As an example, we would like to assign a score to $(e_1, R, e_2) = (\text{Paris}, \text{LOCATED_IN}, \text{France})$ that reflects a high probability.

¹nlp.stanford.edu/projects/snli/, as of February 2019.

Model	Scoring function	Parameters
Unstructured	$\ e_1 - e_2\ _p$	-
TransE	$\ e_1 + w_r - e_2\ _p$	$w_r \in \mathbb{R}^d$
RESCAL	$e_1^T W_r e_2$	$W_r \in \mathbb{R}^{d^2}$
DistMult	$\langle e_1, w_r, e_2 \rangle$	$w_r \in \mathbb{R}^d$
ComplEx	$\text{Re} \langle e_1, w_r, \bar{e}_2 \rangle$	$w_r \in \mathbb{C}^d$

Table 1: Selected relational learning models. Unstructured is from (Bordes et al., 2013a), TransE from (Bordes et al., 2013b), RESCAL from (Nickel et al., 2011), DistMult from (Yang et al., 2015) and (Trouillon et al., 2016). Following the latter, $\langle a, b, c \rangle$ denotes $\sum_k a_k b_k c_k$. $\text{Re}(x)$ is the real part of x , and p is commonly set to 1.

In embedding-based SRL models, entities e_i have vector representations in \mathbb{R}^d and a scoring function reflects truth values of relations. The scoring function should allow for relation-dependent reasoning over the latent space of entities. Scoring functions can have relation-specific parameters, which can be interpreted as relation embeddings. Table 1 presents an overview of a number of state of the art relational models. We can distinguish two families of models: subtractive and multiplicative.

The TransE scoring function is motivated by the idea that translations in latent space can model analogical reasoning and hierarchical relationships. Dense word embeddings trained on tasks related to the distributional hypothesis naturally allow for analogical reasoning with translations without explicit supervision (Mikolov et al., 2013). TransE generalizes the older Unstructured model. We call them subtractive models.

The RESCAL, Distmult, and ComplEx scoring functions can be seen as dot product matching between e_1 and a relation-specific linear transformation of e_2 (Liu et al., 2017). This transformation helps checking whether e_1 matches with some aspects of e_2 . RESCAL allows a full linear mapping $W_r e_2$ but has a high complexity, while Distmult is restricted to a component-wise weighting $w_r \odot e_2$. ComplEx has fewer parameters than RESCAL but still allows for the modeling of asymmetrical relations. As shown in Liu et al. (2017), ComplEx boils down to a restriction of RESCAL where W_r is a block diagonal matrix. These blocks are 2-dimensional, antisymmetric and have equal diagonal terms. Using such a form, even and odd indexes of e 's dimensions play the roles of real and imaginary numbers respectively. The ComplEx model (Trouillon et al., 2016) and its variations

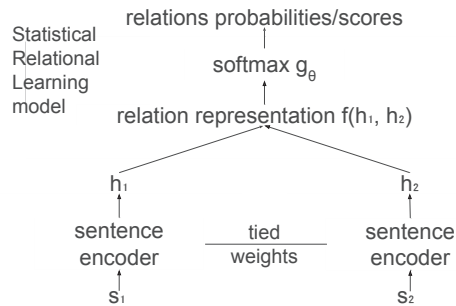


Figure 1: Implicit SRL model in text relation prediction

(Lacroix et al., 2018) yield state of the art performance on knowledge base completion on numerous evaluations.

4 Embeddings composition as SRL models


We claim that several existing models (Conneau et al., 2017; Nie et al., 2017; Baudiš et al., 2016) boil down to SRL models where the sentence embeddings (h_1, h_2) act as entity embeddings (e_1, e_2) . This framework is depicted in figure 1. In this article we focus on sentence embeddings, although our framework can straightforwardly be applied to other levels of language granularity (such as words, clauses, or documents).

Some models (Chen et al., 2017b; Seo et al., 2017; Gong et al., 2018; Radford et al., 2018; Devlin et al., 2018) do not rely on explicit sentence encodings to perform relation prediction. They combine information of input sentences at earlier stages, using conditional encoding or cross-attention. There is however no straightforward way to derive transferable sentence representations in this setting, and so these models are out of the scope of this paper. They sometimes make use of composition functions, so our work could still be relevant to them in some respect.


In this section we will make a link between sentence composition functions and SRL scoring functions, and propose new scoring functions drawing inspiration from SRL.

4.1 Linking composition functions and SRL models

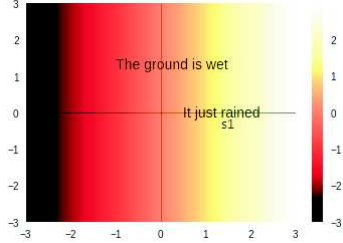
The composition function f_{\odot} from equation 2 followed by a softmax regression yields a score whose analytical form is identical to the Distmult model score described in section 3. Let θ_R denote the softmax weights for relation R . The logit score for the truth of (s_1, R, s_2) is $f(h_1, h_2)\theta_R =$




(a) Score map of $(s_1, R_{to.the.past}, s_2)$ over possible sentences s_2 using Unstructured composition.



(b) Score map of $(s_1, R_{to.the.past}, s_2)$ over possible sentences s_2 using TransE composition.



(c) Score map of $(s_1, R_{entailment}, s_2)$ over possible sentences s_2 using DistMult composition.



(d) Score map of $(s_1, R_{entailment}, s_2)$ over possible sentences s_2 using ComplEx composition.

Figure 2: Possible scoring function values according to different composition functions. s_1 and R are fixed and color brightness reflects likelihood of (s_1, R, s_2) for each position of embedding s_2 . (b) and (d) are respectively more expressive than (a) and (c).

$(h_1 \odot h_2)\theta_R$ which is equal to the Distmult scoring function $\langle h_1, \theta_R, h_2 \rangle$ if h_1, h_2 act as entities embeddings and θ_R as the relation weight w_R .

Similarly, the composition f_- from equation 3 followed by a softmax regression can be seen as an element-wise weighted score of Unstructured (both are equal if softmax weights are all unitary).

Thus, f_{\odot_-} from 4 (with softmax regression) can be seen as a weighted ensemble of Unstructured and Distmult. These two models are respectively outperformed by TransE and ComplEx on knowledge base link prediction by a large margin (Trouillon et al., 2016; Bordes et al., 2013a). We therefore propose to change the Unstructured and Distmult in f_{\odot_-} such that they match their respective state of the art variations in the following sections. We will also show the implications of these refinements.

4.2 Casting TransE as a composition

Simply replacing $|h_2 - h_1|$ with

$$f_t(h_1, h_2) = |h_2 - h_1 + t|, \text{ where } t \in \mathbb{R}^d \quad (6)$$

would make the model analogous to TransE. t is learned and is shared by all relations. A relation-specific translation t_R could be used but it would make f relation-specific. Instead, here, each dimension of $f_t(h_1, h_2)$ can be weighted according

to a given relation. Non-zero t makes f_t asymmetrical and also yields features that allow for the checking of an analogy between s_1 and s_2 . Sentence embeddings often rely on pre-trained word embeddings which have demonstrated strong capabilities for analogical reasoning. Some analogies, such as *part-whole*, are computable with off-the-shelf word embeddings (Chen et al., 2017a) and should be very informative for natural language inference tasks. As an illustration, let us consider an artificial semantic space (depicted in figures 2a and 2b) where we posit that there is a “to the past” translation t so that $h_1 + t$ is the embedding of a sentence s_1 changed to the past tense. Unstructured is not able to leverage this semantic space to correctly score $(s_1, R_{to.the.past}, s_2)$ while TransE is well tailored to provide highest scores for sentences near $h_1 + \hat{t}$ where \hat{t} is an estimation of t that could be learned from examples.

4.3 Casting ComplEx as a composition

Let us partition h dimensions into two equally sized sets \mathcal{R} and \mathcal{I} , e.g. even and odd dimension indices of h . We propose a new function f_C as a way to fit the ComplEx scoring function into a composition function.

$$f_C(h_1, h_2) = [h_1^{\mathcal{R}} \odot h_2^{\mathcal{R}} + h_1^{\mathcal{I}} \odot h_2^{\mathcal{I}}, \quad (7) \\ h_1^{\mathcal{R}} \odot h_2^{\mathcal{I}} - h_1^{\mathcal{I}} \odot h_2^{\mathcal{R}}]$$

$f_{\mathbb{C}}(h_1, h_2)$ multiplied by softmax weights θ_r is equivalent to the ComplEx scoring function $\text{Re} \langle h_1, \theta_r, \overline{h_2} \rangle$. The first half of θ_r weights corresponds to the real part of ComplEx relation weights while the last half corresponds to the imaginary part.

$f_{\mathbb{C}}$ is to the ComplEx scoring function what f_{\odot} is to the DistMult scoring function. Intuitively, ComplEx is a minimal way to model interactions between distinct latent dimensions while Distmult only allows for identical dimensions to interact.

Let us consider a new artificial semantic space (shown in figures 2c and 2d) where the first dimension is high when a sentence means that it just rained, and the second dimension is high when the ground is wet. Over this semantic space, Distmult is only able to detect entailment for paraphrases whereas ComplEx is also able to naturally model that (“it just rained”, $R_{\text{entailment}}$, “the ground is wet”) should be high while its converse should not.

We also propose two more general versions of $f_{\mathbb{C}}$:

$$f_{\mathbb{C}^{\alpha}}(h_1, h_2) = [h_1^{\mathcal{R}} \odot h_2^{\mathcal{R}}, h_1^{\mathcal{I}} \odot h_2^{\mathcal{I}}, h_1^{\mathcal{R}} \odot h_2^{\mathcal{I}} - h_1^{\mathcal{I}} \odot h_2^{\mathcal{R}}] \quad (8)$$

$$f_{\mathbb{C}^{\beta}}(h_1, h_2) = [h_1^{\mathcal{R}} \odot h_2^{\mathcal{R}}, h_1^{\mathcal{I}} \odot h_2^{\mathcal{I}}, h_1^{\mathcal{R}} \odot h_2^{\mathcal{I}}, h_1^{\mathcal{I}} \odot h_2^{\mathcal{R}}] \quad (9)$$

$f_{\mathbb{C}^{\alpha}}$ can be seen as Distmult concatenated with the asymmetrical part of ComplEx and $f_{\mathbb{C}^{\beta}}$ can be seen as RESCAL with unconstrained block diagonal relation matrices.

5 On the evaluation of relational models

The SentEval framework (Conneau et al., 2017) provides a general evaluation for transferable sentence representations, with open source evaluation code. One only needs to specify a sentence encoder function, and the framework performs classification tasks or relation prediction tasks using cross-validated logistic regression on embeddings or composed sentence embeddings. Tasks include sentiment analysis, entailment, textual similarity, textual relatedness, and paraphrase detection. These tasks are a rich way to train or evaluate sentence representations since in a triple (s_1, R, s_2) , we can see (R, s_2) as a label for s_1 (Baudiš et al., 2016). Unfortunately, the relational tasks hard-code the composition function

from equation 4. From our previous analysis, we believe this composition function favors the use of contextual/lexical similarity rather than high-level reasoning and can penalize representations based on more semantic aspects. This bias could harm research since semantic representation is an important next step for sentence embedding. Training/evaluation datasets are also arguably flawed with respect to relational aspects since several recent studies (Dasgupta et al., 2018; Poliak et al., 2018; Gururangan et al., 2018; Glockner et al., 2018) show that InferSent, despite being state of the art on SentEval evaluation tasks, has poor performance when dealing with asymmetrical tasks and non-additive composition of words. In addition to providing new ways of training sentence encoders, we will also extend the SentEval evaluation framework with a more expressive composition function when dealing with relational transfer tasks, which improves results even when the sentence encoder was not trained with it.

6 Experiments

Our goal is to show that transferable sentence representation learning and relation prediction tasks can be improved when our expressive compositions are used instead of the composition from equation 4. We train our relational model adaptations on two relation prediction base tasks (\mathcal{T}), one supervised ($\mathcal{T} = NLI$) and one unsupervised ($\mathcal{T} = Disc$) described below, and evaluate sentence/relation representations on base and transfer tasks using the SentEval framework in order to quantify the generalization capabilities of our models. Since we use minor modifications of InferSent and SentEval, our experiments are easily reproducible.

6.1 Training tasks

Natural language inference ($\mathcal{T} = NLI$)’s goal is to predict whether the relation between two sentences (premise and hypothesis) is *Entailment*, *Contradiction* or *Neutral*. We use the combination of SNLI dataset (Bowman et al., 2015) and MNLI dataset (Williams et al., 2018). We call *AllNLI* the resulting dataset of 1M examples. Conneau et al. (2017) claim that NLI data allows universal sentence representation learning. They used the $f_{\odot,-}$ composition function with concatenated sentence representations in order to train their *Infersent* model.

name	N	task	C	representation(s) used
MR	11k	sentiment (movies)	2	h_1
SUBJ	10k	subjectivity/objectivity	2	h_1
MPQA	11k	opinion polarity	2	h_1
TREC	6k	question-type	6	h_1
SICK _s ^m	10k	NLI	3	$f_{m,s}(h_1, h_2)$
MRPC _s ^m	4k	paraphrase detection	2	$(f_{m,s}(h_1, h_2) + (f_{m,s}(h_2, h_1)))/2$
PDTB _s ^m	17k	discursive relation	5	$f_{m,s}(h_1, h_2)$
STS14	4.5k	similarity	-	$\cos(h_1, h_2)$

Table 2: Transfer evaluation tasks. N = number of training examples; C = number of classes if applicable. h_1, h_2 are sentence representations, $f_{m,s}$ a composition function from section 4.

We also train on the prediction of discourse connectives between sentences/clauses ($\mathcal{T} = \text{Disc}$). Discourse connectives make discourse relations between sentences explicit. In the sentence *I live in Paris but I'm often elsewhere*, the word *but* highlights that there is a contrast between the two clauses it connects. We use Malmi et al.'s (2018) dataset of selected 400k instances with 20 discourse connectives (e.g. *however, for example*) with the provided train/dev/test split. This dataset has no other supervision than the list of 20 connectives. Nie et al. (2017) used $f_{\odot,-}$ concatenated with the sum of sentence representations to train their model, *DisSent*, on a similar task and showed that their encoder was general enough to perform well on SentEval tasks. They use a dataset that is, at the time of writing, not publicly available.

6.2 Evaluation tasks

Table 2 provides an overview of different transfer tasks that will be used for evaluation. We added another relation prediction task, the PDTB coarse-grained implicit discourse relation task, to SentEval. This task involves predicting a discursive link between two sentences among {Comparison, Contingency, Entity based coherence, Expansion, Temporal}. We followed the setup of Pitler et al. (2009), without sampling negative examples in training. MRPC, PDTB and SICK will be tested with two composition functions: besides SentEval composition $f_{\odot,-}$, we will use $f_{\mathbb{C}^\beta,-}$ for transfer learning evaluation, since it has the most general multiplicative interaction and it does not penalize models that do not learn a translation. For all tasks except STS14, a cross-validated logistic regression is used on the sentence or relation representation. The evaluation of the STS14 task relies on Pearson or Spearman correlation between cosine similarity and the target. We force the composition function to be symmetrical on the MRPC

task since paraphrase detection should be invariant to permutation of input sentences.

6.3 Setup

We want to compare the different instances of f . We follow the setup of Infersent (Conneau et al., 2017): we learn to encode sentences into h with a bi-directional LSTM using element-wise max pooling over time. The dimension size of h is 4096. Word embeddings are fixed GloVe with 300 dimensions, trained on Common Crawl 840B.² Optimization is done with SGD and decreasing learning rate until convergence.

The only difference with regard to Infersent is the composition. Sentences are composed with six different compositions for training according to the following template:

$$f_{m,s,1,2}(h_1, h_2) = [f_m(h_1, h_2), f_s(h_1, h_2), h_1, h_2] \quad (10)$$

f_s (subtractive interaction) is in $\{f_-, f_t\}$, f_m (multiplicative interaction) is in $\{f_{\odot}, f_{\mathbb{C}^\alpha}, f_{\mathbb{C}^\beta}\}$. We do not consider $f_{\mathbb{C}}$ since it yielded inferior results in our early experiments using NLI and SentEval development sets.

$f_{m,s,1,2}(h_1, h_2)$ is fed directly to a softmax regression. Note that Infersent uses a multi-layer perceptron before the softmax, but uses only linear activations, so $f_{\odot,-,1,2}(h_1, h_2)$ is analytically equivalent to Infersent when $\mathcal{T} = \text{NLI}$.

6.4 Results

Having run several experiments with different initializations, the standard deviations between them do not seem to be negligible. We decided to take these into account when reporting scores, contrary to previous work (Kiros et al., 2015; Conneau et al., 2017): we average the scores of 6 distinct runs for each task and use standard deviations

²<https://nlp.stanford.edu/projects/glove/>

Models trained on natural language inference ($\mathcal{T} = NLI$)										
m,s	MR	SUBJ	MPQA	TREC	MRPC $_{\ominus}$	PDTB $_{\ominus}$	SICK $_{\ominus}$	STS14	\mathcal{T}	AVG
$\odot, -$	81.2	92.7	90.4	89.6	76.1	46.7	86.6	69.5	84.2	79.1
$\alpha, -$	81.4	92.8	90.5	89.6	75.4	46.6	86.7	69.5	84.3	79.1
$\beta, -$	81.2	92.6	90.5	89.6	76	46.5	86.6	69.5	84.2	79.1
\odot, t	81.1	92.7	90.5	89.7	76.5	46.4	86.5	70.0	84.8	79.2
α, t	81.3	92.6	90.6	89.2	76.2	47.2	86.5	70.0	84.6	79.2
β, t	81.2	92.7	90.4	88.5	75.8	47.3	86.8	69.8	84.2	79.1

Table 3: SentEval and base task evaluation results for the models trained on natural language inference ($\mathcal{T} = NLI$); AllNLI is used for training. All scores are accuracy percentages, except STS14, which is Pearson correlation percentage. AVG denotes the average of the SentEval scores.

Models trained on discourse connective prediction ($\mathcal{T} = Disc$)										
m,s	MR	SUBJ	MPQA	TREC	MRPC $_{\ominus}$	PDTB $_{\ominus}$	SICK $_{\ominus}$	STS14	\mathcal{T}	AVG
$\odot, -$	80.4	92.7	90.2	89.5	74.5	47.3	83.2	57.9	35.7	77
$\alpha, -$	80.4	92.9	90.2	90.2	75	47.9	83.3	57.8	35.9	77.2
$\beta, -$	80.2	92.8	90.2	88.4	74.9	47.5	82.9	57.7	35.9	76.8
\odot, t	80.2	92.8	90.2	90.4	74.6	48.5	83.4	58.6	36.1	77.3
α, t	80.2	92.9	90.3	90.3	75.1	47.8	83.2	58.3	36.1	77.3
β, t	80.2	92.8	90.3	89.7	74.4	47.9	83.7	58.2	35.7	77.2

Table 4: SentEval and base task evaluation results for the models trained on discourse connective prediction ($\mathcal{T} = Disc$). All scores are accuracy percentages, except STS14, which is Pearson correlation percentage. AVG denotes the average of the SentEval scores.

under normality assumption to compute significance. Table 3 shows model scores for $\mathcal{T} = NLI$, while Table 4 shows scores for $\mathcal{T} = Disc$. For comparison, Table 5 shows a number of important models from previous work. Finally, in Table 6, we present results for sentence relation tasks that use an alternative composition function ($f_{\mathbb{C}^{\beta}, -}$) instead of the standard composition function used in SentEval.

For sentence representation learning, the baseline, $f_{\odot, -}$ composition already performs rather well, being on par with the InferSent scores of the original paper, as would be expected. However, macro-averaging all accuracies, it is the second worst performing model. $f_{\mathbb{C}^{\alpha}, t, 1, 2}$ is the best performing model, and all three best models use the translation ($s = t$). On relational transfer tasks, training with $f_{\mathbb{C}^{\alpha}, t, 1, 2}$ and using complex \mathbb{C}^{β} for transfer (Table 6) always outperforms the baseline ($f_{\odot, -, 1, 2}$ with $\odot, -$ composition in Tables 3 and 4). Averaging accuracies of those transfer tasks, this result is significant for both training tasks at level $p < 0.05$ (using Bonferroni correction accounting for the 5 comparisons). On base tasks and the average of non-relational transfer tasks (MR, MPQA, SUBJ, TREC), our proposed compositions are on average slightly better than $f_{\odot, -, 1, 2}$. Representations learned with our proposed compositions can

still be compared with simple cosine similarity: all three methods using the translational composition ($s = t$) very significantly outperform the baseline (significant at level $p < 0.01$ with Bonferroni correction) on STS14 for $\mathcal{T} = NLI$. Thus, we believe $f_{\mathbb{C}^{\alpha}, t, 1, 2}$ has more robust results and could be a better default choice than $f_{\odot, -, 1, 2}$ as composition for representation learning.³

Additionally, using \mathbb{C}^{β} (Table 6) instead of \odot (Tables 3 and 4) for transfer learning in relational transfer tasks (PDTB, MRPC, SICK) yields a significant improvement on average, even when $m = \odot$ was used for training ($p < 0.001$). Therefore, we believe $f_{\mathbb{C}^{\beta}, -}$ is an interesting composition for inference or evaluation of models regardless of how they were trained.

7 Related work

There are numerous interactions between SRL and NLP. We believe that our framework merges two specific lines of work: relation prediction and modeling textual relational tasks.

Some previous NLP work focused on composition functions for relation prediction between

³Note that our compositions are also beneficial with regard to convergence speed: on average, each of our proposed compositions needed less epochs to converge than the baseline $f_{\odot, -, 1, 2}$, for both training tasks.

Comparison models									
model	MR	SUBJ	MPQA	TREC	MRPC [⊖]	PDTB [⊖]	SICK [⊖]	STS14	AVG
InferSent	81.1	92.4	90.2	88.2	76.2	46.7-	86.3	70	78.9
SkipT	76.5	93.6	87.1	92.2	73	-	82.3	29	-
BoW	77.2	91.2	87.9	83	72.2	43.9	78.4	54.6	73.6

Table 5: Comparison models from previous work. InferSent represents the original results from Conneau et al. (2017), SkipT is SkipThought from Kiros et al. (2015), and BoW is our re-evaluation of GloVe Bag of Words from Conneau et al. (2017). AVG denotes the average of the SentEval scores..

m,s	$\mathcal{T} = Disc$				$\mathcal{T} = NLI$			
	MRPC ^β	PDTB ^β	SICK ^β	AVG	MRPC ^β	PDTB ^β	SICK ^β	AVG
⊖, −	74.8	48.2	83.6	68.9	76.2	47.2	86.9	70.1
α, −	74.9	49.3	83.8	69.3	75.9	47.1	86.9	70
β, −	75	48.8	83.4	69.1	75.8	47	87	69.9
⊖, t	74.9	48.7	83.6	69.1	76.2	47.8	86.8	70.3
α, t	75.2	48.6	83.5	69.1	76.2	47.6	87.3	70.4
β, t	74.6	48.9	83.9	69.1	76.2	47.8	87	70.3

Table 6: Results for sentence relation tasks using an alternative composition function ($f_{C^\beta, -}$) during evaluation. AVG denotes the average of the three tasks.

text fragments, even though they ignored SRL and only dealt with word units. Word2vec (Mikolov et al., 2013) has sparked a great interest for this task with word analogies in the latent space. Levy & Goldberg (2014) explored different scoring functions between words, notably for analogies. Hypernymy relations were also studied, by Chang et al. (2018) and Fu et al. (2014). Levy et al. (2015) proposed tailored scoring functions. Even the skipgram model (Mikolov et al., 2013) can be formulated as finding relations between context and target words. We did not empirically explore textual relational learning at the word level, but we believe that it would fit in our framework, and could be tested in future studies. Numerous approaches (Chen et al., 2017b; Seok et al., 2016; Gong et al., 2018; Joshi et al., 2019) were proposed to predict inference relations between sentences, but don’t explicitly use sentence embeddings. Instead, they encode sentences jointly, possibly with the help of previously cited word compositions, therefore it would also be interesting to try applying our techniques within their framework.

Some modeling aspects of textual relational learning have been formally investigated by Baudiš et al. (2016). They noticed the genericity of relational problems and explored multi-task and transfer learning on relational tasks. Their work is complementary to ours since their framework unifies tasks while ours unifies composition func-

tions. Subsequent approaches use relational tasks for training and evaluation on specific datasets (Conneau et al., 2017; Nie et al., 2017).

8 Conclusion

We have demonstrated that a number of existing models used for textual relational learning rely on composition functions that are already used in Statistical Relational Learning. By taking into account previous insights from SRL, we proposed new composition functions and evaluated them. These composition functions are all simple to implement and we hope that it will become standard to try them on relational problems. Larger scale data might leverage these more expressive compositions, as well as more compositional, asymmetric, and arguably more realistic datasets (Dasgupta et al., 2018; Gururangan et al., 2018). Finally, our compositions can also be helpful to improve interpretability of embeddings, since they can help measure relation prediction asymmetry. Analogies through translations helped interpreting word embeddings, and perhaps analyzing our learned t translation could help interpreting sentence embeddings.

References

- Isaac K E Ampomah, Seong-bae Park, and Sang-jo Lee. A Sentence-to-Sentence Relation Network for Recognizing Textual Entailment. *World Academy of Science, Engineering and Technology International Journal of Computer and Information Engineering*, 10(12):1955–1958, 2016.
- Petr Baudiš, Jan Pichl, Tomáš Vyskočil, and Jan Šedivý. Sentence Pair Scoring: Towards Unified Framework for Text Comprehension. 2016. URL <http://arxiv.org/abs/1603.06127>.
- Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. A Semantic Matching Energy Function for Learning with Multi-relational Data. *Machine Learning*, 2013a. ISSN 0885-6125. doi: 10.1007/s10994-013-5363-6. URL <http://arxiv.org/abs/1301.3485>.
- Antoine Bordes, Nicolas Usunier, Jason Weston, and Oksana Yakhnenko. Translating Embeddings for Modeling Multi-Relational Data. *Advances in NIPS*, 26:2787–2795, 2013b. ISSN 10495258. doi: 10.1007/s13398-014-0173-7.2.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. A large annotated corpus for learning natural language inference. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, Lisbon, Portugal, 17-21 September 2015*, (September):632–642, 2015. ISSN 9781941643327.
- Haw-Shiuan Chang, Ziyun Wang, Luke Vilnis, and Andrew McCallum. Distributional inclusion vector embedding for unsupervised hypernymy detection. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 485–495, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1045. URL <https://www.aclweb.org/anthology/N18-1045>.
- Dawn Chen, Joshua C. Peterson, and Tom Griffiths. Evaluating vector-space models of analogy. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society, CogSci 2017, London, UK, 16-29 July 2017*, 2017a. URL <https://mindmodeling.org/cogsci2017/papers/0340/index.html>.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. Enhanced lstm for natural language inference. In Regina Barzilay and Min-Yen Kan (eds.), *ACL (1)*, pp. 1657–1668. Association for Computational Linguistics, 2017b. ISBN 978-1-945626-75-3. URL <http://dblp.uni-trier.de/db/conf/acl/acl2017-1.html#ChenZLWJI17>.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. Supervised Learning of Universal Sentence Representations from Natural Language Inference Data. *Emnlp*, 2017.
- Ishita Dasgupta, Demi Guo, Andreas Stuhlmüller, Samuel J. Gershman, and Noah D. Goodman. Evaluating Compositionality in Sentence Embeddings. (2011), 2018. URL <http://arxiv.org/abs/1802.04302>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Ruiji Fu, Jiang Guo, Bing Qin, Wanxiang Che, Haifeng Wang, and Ting Liu. Learning semantic hierarchies via word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1199–1209, Baltimore, Maryland, June 2014. Association for Computational Linguistics. doi: 10.3115/v1/P14-1113. URL <https://www.aclweb.org/anthology/P14-1113>.
- Lise Getoor and Ben Taskar. *Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2007. ISBN 0262072882.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. Breaking NLI Systems with Sentences that Require Simple Lexical Inferences. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Short Papers)*, (3):1–6, 2018. URL <http://arxiv.org/abs/1805.02266>.
- Yichen Gong, Heng Luo, and Jian Zhang. Natural language inference over interaction space. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rldHXnH6->.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pp. 107–112, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-2017. URL <https://www.aclweb.org/anthology/N18-2017>.
- Sohail Hooda and Leila Kosseim. Argument labeling of explicit discourse relations using lstm neural networks. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pp. 309–315, Varna, Bulgaria, September 2017. INCOMA Ltd. doi: 10.26615/978-954-452-049-6_042. URL https://doi.org/10.26615/978-954-452-049-6_042.

- Yacine Jernite, Samuel R. Bowman, and David Sontag. Discourse-Based Objectives for Fast Unsupervised Sentence Representation Learning. 2017. URL <http://arxiv.org/abs/1705.00557>.
- Mandar Joshi, Eunsol Choi, Omer Levy, Daniel S. Weld, and Luke Zettlemoyer. pair2vec: Compositional word-pair embeddings for cross-sentence inference. In *NAACL*, 2019. URL <http://arxiv.org/abs/1810.08854>.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. In *Advances in neural information processing systems*, pp. 3294–3302, 2015.
- Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. Canonical tensor decomposition for knowledge base completion. In *ICML*, 2018.
- Omer Levy and Yoav Goldberg. Linguistic Regularities in Sparse and Explicit Word Representations. *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pp. 171–180, 2014. doi: 10.3115/v1/W14-1618. URL <http://aclweb.org/anthology/W14-1618>.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. Do Supervised Distributional Methods Really Learn Lexical Inference Relations? *Naacl-2015*, pp. 970–976, 2015. URL <http://www.aclweb.org/anthology/N15/N15-1098.pdf>.
- Hanxiao Liu, Yuexin Wu, and Yiming Yang. Analytical Inference for Multi-Relational Embeddings. *Icml*, 2017. ISSN 1938-7228. URL <http://arxiv.org/abs/1705.02426>.
- Lajanugen Logeswaran, Honglak Lee, and Dragomir Radev. Sentence Ordering using Recurrent Neural Networks. pp. 1–15, 2016. URL <http://arxiv.org/abs/1611.02654>.
- Lajanugen Logeswaran, Honglak Lee, and Dragomir R. Radev. Sentence ordering and coherence modeling using recurrent neural networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pp. 5285–5292, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17011>.
- Eric Malmi, Daniele Pighin, Sebastian Krause, and Mikhail Kozhevnikov. Automatic prediction of discourse connectives. In *Proceedings of the 11th Language Resources and Evaluation Conference*, Miyazaki, Japan, May 2018. European Language Resource Association. URL <https://www.aclweb.org/anthology/L18-1260>.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed Representations of Words and Phrases and their Compositionality. *Nips*, pp. 1–9, 2013. ISSN 10495258. doi: 10.1162/jmlr.2003.3.4-5.951.
- Lili Mou, Rui Men, Ge Li, Yan Xu, Lu Zhang, Rui Yan, and Zhi Jin. Natural language inference by tree-based convolution and heuristic matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 130–136, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-2022. URL <https://www.aclweb.org/anthology/P16-2022>.
- Jonas Mueller and Aditya Thyagarajan. Siamese Recurrent Architectures for Learning Sentence Similarity. In *AAAI*, pp. 2786–2792, 2016.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A Three-Way Model for Collective Learning on Multi-Relational Data. *Icml*, pp. 809–816, 2011.
- Allen Nie, Erin D. Bennett, and Noah D. Goodman. DisSent: Sentence Representation Learning from Explicit Discourse Relations. 2017. URL <http://arxiv.org/abs/1710.04334>.
- Emily Pitler, Annie Louis, and Ani Nenkova. Automatic sense prediction for implicit discourse relations in text. *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP Volume 2 ACLIJCNLP 09*, 2(August):683–691, 2009. doi: 10.3115/1690219.1690241. URL <http://www.aclweb.org/anthology/P/P09/P09-1077>.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. Hypothesis Only Baselines in Natural Language Inference. *Proceedings of the 7th Joint Conference on Lexical and Computational Semantics*, (1):180–191, 2018.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. <https://www.cs.ubc.ca/~amuham01/LING530/papers/radford2018improving.pdf>, 2018.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017. URL <https://openreview.net/forum?id=HJ0UKP9ge>.
- Miran Seok, Hye-Jeong Song, Chan-Young Park, Jong-Dae Kim, and Yu-Seop Kim. Named Entity Recognition using Word Embedding as a Feature 1. *International Journal of Software Engineering and Its Applications*, 10(2):93–104, 2016.

ISSN 1738-9984. doi: 10.14257/ijseia.2016.10.2.08. URL <http://dx.doi.org/10.14257/ijseia.2016.10.2.08>.

Richard Socher, Danqi Chen, Christopher Manning, Danqi Chen, and Andrew Ng. Reasoning With Neural Tensor Networks for Knowledge Base Completion. In *Neural Information Processing Systems (2013)*, pp. 926–934, 2013. URL <https://nlp.stanford.edu/pubs/SocherChenManningNg{ }NIPS2013.pdf>.

Kai Sheng Tai, Richard Socher, and Christopher D. Manning. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pp. 1556–1566, 2015. URL <http://aclweb.org/anthology/P/P15/P15-1150.pdf>.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex Embeddings for Simple Link Prediction. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48, 2016. ISBN 9781510829008. URL <http://arxiv.org/abs/1606.06357>.

Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1101. URL <https://www.aclweb.org/anthology/N18-1101>.

Min Chul Yang, Do Gil Lee, So Young Park, and Hae Chang Rim. Knowledge-based question answering using the semantic embedding space. *Expert Systems with Applications*, 42(23):9086–9104, 2015. ISSN 09574174. doi: 10.1016/j.eswa.2015.07.009. URL <http://dx.doi.org/10.1016/j.eswa.2015.07.009>.