

---

# Convolutional 2D Knowledge Graph Embeddings

---

**Tim Dettmers\***

Università della Svizzera italiana  
tim.dettmers@gmail.com

**Pasquale Minervini Pontus Stenetorp Sebastian Riedel**

University College London  
{p.minervini, p.stenetorp, s.riedel}@ucl.ac.uk

## Abstract

In this work, we introduce a convolutional neural network model, ConvE, for the task of link prediction. ConvE applies 2D convolution directly on embeddings, thus inducing spatial structure in embedding space. To scale to large knowledge graphs and prevent overfitting due to over-parametrization, previous work seeks to reduce parameters by performing simple transformations in embedding space. We take inspiration from computer vision, where convolution is able to learn multiple layers of non-linear features while reducing the number of parameters through weight sharing. Applied naively, convolutional models for link prediction are computationally costly. However, by predicting all links simultaneously we improve test time performance by more than 300x on FB15k. We report state-of-the-art results for numerous previously introduced link prediction benchmarks, including the well-established FB15k and WN18 datasets. Previous work noted that these two datasets contain many reversible triples, but the severity of this issue was not quantified. To investigate this, we design a simple model that uses a single rule which reverses relations and achieves state-of-the-art results. We introduce WN18RR, a subset of WN18 which was constructed the same way as the previously proposed FB15k-237, to alleviate this problem and report results for our own and previously proposed models for all datasets. Analysis of our convolutional model suggests that it is particularly good at modelling nodes with high indegree and nodes with high PageRank and that 2D convolution applied on embeddings seems to induce contrasting pixel-level structures.

## 1 Introduction

Knowledge graphs are databases of triples, where facts are represented in the form of relationships (edges) between entities (nodes). While knowledge graphs have important applications in search, analytics, recommendation, and data integration they suffer from incompleteness, that is, missing links in the graph. For example, in Freebase and DBpedia more than 66% of the person entries are missing a birthplace [6, 15]. Link prediction is the task to predict missing links in such a graph.

One focus of previous link prediction work lies in the number of model parameters. Knowledge graphs can be large and as such link predictors should scale in a manageable way with respect to both the number of parameters and computational costs to be applicable in real-world scenarios. It has also been shown that more complex models suffer from over-parameterisation and are thus prone to overfitting [22]. Due to these limitations, link prediction models are often composed of simple operations, like dot products and matrix multiplications, over an embedding space and use a limited number of parameters.

---

\*This work was conducted during a research visit to University College London.

In computer vision, convolutional layers are used to overcome these problems by using parameter sharing to learn multiple layers of non-linear features which are increasingly abstract and as such more expressive than shallower models [14].

Due to a range of regularisation techniques [9, 28, 30], convolutional architectures have little or no problems with overfitting even for networks with a depth close to a hundreds layers [29, 8]. Convolutional layers are highly parameter efficient, for example convolutional architectures have about 90% of their parameters in the fully connected hidden layers which compose features, while having only 10% of the parameters in convolutional layers which extract features [14]. As such, convolutional layers constitute highly parameter efficient feature extractors which can model rich non-linear interactions and generalise effectively.

In this work we propose a neural link predictor, ConvE, that uses 2D convolution over embeddings to predict new links in knowledge graphs. While the bulk of our model’s parameters is still in the relation and entity embeddings like in other link predictors, by using convolution as a weight sharing mechanism, we only use an additional 72 parameters to extract an extra layer of non-linear features which are then projected back to embedding space for scoring. We thus have a highly parameter efficient, scalable architecture, which generalises well and uses few additional parameters compared to other commonly used link prediction models. We make the following contributions:

- We introduce ConvE, a simple model that uses 2D convolutions over embeddings and achieves state-of-the-art performance on the WN18, FB15k, YAGO3-10 and Countries datasets.
- Using a 1-N approach for link prediction, where we predict scores for all possible links simultaneously, we speed up evaluation by several orders of magnitude. While useful for our model to get more predictions from each expensive convolution operation, this approach can also help speed up other link prediction models.
- Previous work by Toutanova and Chen [31] noted that FB15k and WN18 contain many redundant, reversible relations, but they did not investigate the severity of this problem. We demonstrate the severity by designing a simple model which is based on a reversal rule that achieves state-of-the-art results on WN18 and FB15k, suggesting that successful models might learn this rule rather than to model the knowledge graph itself. We propose a new version of WN18, which follows the construction procedure of FB15k-237 to alleviate this problem.
- We show that the performance of our model compared to a shallow model, DistMult, increases when the network and test set involves nodes with high indegree or high PageRank. In particular, we show that the performance of our model increases relative to DistMult with increasing mean PageRank of nodes in the test set ( $r = 0.83$ ). This suggests that our model is better at modelling nodes with high indegree.

## 2 Related Work

Several neural link prediction models have been proposed in the literature, such as the Translating Embeddings model (TransE) [1], the Bilinear Diagonal model (DistMult) [34] and its extension in the complex space (ComplEx) [33] – we refer to Nickel et al. [22] for a recent survey on such models. The neural link prediction model that is most closely related to this work is most likely the Holographic Embeddings model (HolE) [23], which uses cross-correlation – the inverse of circular convolution – for matching entity embeddings; it is inspired by holographic models of associative memory. However, HolE does not learn multiple layers of non-linear features and is thus theoretically less expressive than our model.

To the best of our knowledge, our model is the first model to use 2D convolutional layers when defining neural link prediction models for link prediction. *Graph Convolutional Networks* (GCNs) [7, 5, 13] are a related line of research, where the convolution operator is generalised to use locality information in graphs. However, the GCN framework is limited to undirected graphs while knowledge graphs are naturally directed, and suffers from potentially prohibitive memory requirements [13]. Relational GCNs (R-GCNs) [26] are a generalisation of GCNs developed for dealing with highly multi-relational data such as knowledge graphs – we include them in our experimental evaluations.

Table 1: Scoring functions  $\psi_r(\mathbf{e}_s, \mathbf{e}_o)$  from neural link predictors in the literature, their relation-dependent parameters and space complexity.

Model	Score $\psi_r(\mathbf{e}_s, \mathbf{e}_o)$	Relation parameters	Space complexity
RESCAL [21]	$\mathbf{e}_s^T \mathbf{W}_r \mathbf{e}_o$	$\mathbf{W} \in \mathbb{R}^{k \times k}$	$\mathcal{O}(n_e k + n_r k^2)$
SE [3]	$\ \mathbf{W}_r^L \mathbf{e}_s - \mathbf{W}_r^R \mathbf{e}_o\ _p$	$\mathbf{W}_r^L, \mathbf{W}_r^R \in \mathbb{R}^{k \times k}$	$\mathcal{O}(n_e k + n_r k^2)$
TransE [1]	$\ \mathbf{e}_s + \mathbf{r}_r - \mathbf{e}_o\ _p$	$\mathbf{r}_r \in \mathbb{R}^k$	$\mathcal{O}(n_e k + n_r k)$
DistMult [34]	$\langle \mathbf{e}_s, \mathbf{r}_r, \mathbf{e}_o \rangle$	$\mathbf{r}_r \in \mathbb{R}^k$	$\mathcal{O}(n_e k + n_r k)$
ComplEx [33]	$\langle \mathbf{e}_s, \mathbf{r}_r, \mathbf{e}_o \rangle$	$\mathbf{r}_r \in \mathbb{C}^k$	$\mathcal{O}(n_e k + n_r k)$
ConvE	$f(\text{vec}(f([\overline{\mathbf{e}_s}; \overline{\mathbf{r}_r}] * \omega)) \mathbf{W}) \mathbf{e}_o$	$\mathbf{r}_r \in \mathbb{R}^k$	$\mathcal{O}(n_e k + n_r k')$

Several CNN-based models have been proposed in natural language processing (NLP) for solving a variety of tasks, including semantic parsing [35], sentence classification [11], search query retrieval [27], sentence modelling [10], as well as traditional NLP tasks [4]. However, most work in NLP uses 1D-convolutions, that is convolutions which operate over a temporal sequence of embeddings, for example a sequence of words in embedding space. In this work, we use 2D-convolutions which operate on a spatial level directly on embeddings. As we show later in section 7.2, this induces pixel-level spatial structure in our embeddings.

Using 2D convolutions has one major advantage for interactions between embeddings: Consider for example the case where we concatenate two rows of 1D embeddings — that is lining them up — a 1D convolution will be able to extract features of the interaction between these two embeddings at the concatenation point; if we concatenate two rows of 2D embeddings — that is stacking them — a 2D convolution will be able to extract features of interactions over the entire concatenation line. Thus 2D convolution is able to extract more features of interactions between two embeddings compared to 1D convolution.

### 3 Background

A *knowledge graph*  $\mathcal{G} \triangleq \{(s, r, o)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$  can be formalised as a set of triples (facts), each consisting of a relationship  $r \in \mathcal{R}$  and two entities  $s, o \in \mathcal{E}$ , referred to as the *subject* and *object* of the triple. Each triple  $(s, r, o)$  denotes a relationship of type  $r$  between the entities  $s$  and  $o$ .

The *link prediction* problem can be formalised as a pointwise learning to rank problem, where the objective is learning a scoring function  $\psi : \mathcal{E} \times \mathcal{R} \times \mathcal{E} \mapsto \mathbb{R}$ . Given an input triple  $x = (s, r, o)$ , its score  $\psi(x) \in \mathbb{R}$  is proportional to the likelihood that the fact encoded by  $x$  is true.

In our case, the score of a relationships is defined as a deep convolutional network [16].

#### Neural Link Predictors

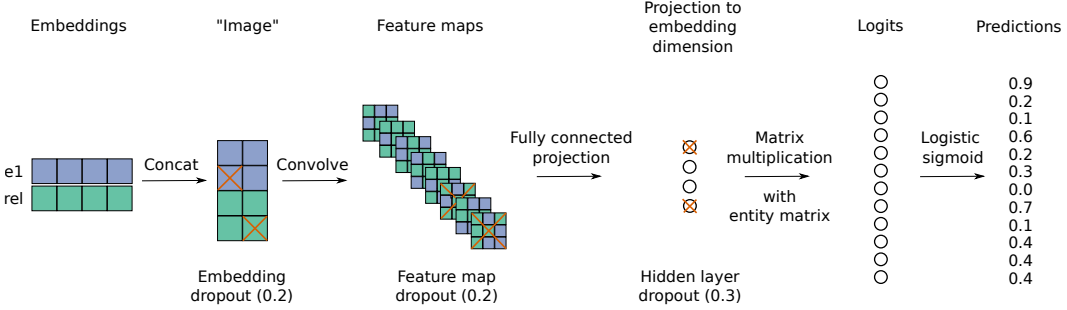
Neural link prediction models [22] can be seen as multi-layer neural networks composed by an *encoding component* and a *scoring component*. Given an input triple  $(s, r, o)$ , the encoding component maps entities  $s, o \in \mathcal{E}$  to their distributed embedding representations  $\mathbf{e}_s, \mathbf{e}_o \in \mathbb{R}^k$ . In the scoring component, the two entity embeddings  $\mathbf{e}_s$  and  $\mathbf{e}_o$  are scored by a function  $\psi_r$ . The score of  $(s, r, o)$  is defined as  $\psi(s, r, o) \triangleq \psi_r(\mathbf{e}_s, \mathbf{e}_o) \in \mathbb{R}$ .

In Table 1 we summarise the scoring function of some link prediction models from the literature. The vectors  $\mathbf{e}_s$  and  $\mathbf{e}_o$  denote the subject and object embedding, where  $\mathbf{e}_s, \mathbf{e}_o \in \mathbb{C}^k$  in ComplEx and  $\mathbf{e}_s, \mathbf{e}_o \in \mathbb{R}^k$  in all other models, and  $\langle x, y, z \rangle \triangleq \sum_i x_i y_i z_i$  denotes the tri-linear dot product;  $*$  denotes the convolution operator;  $f$  denotes a non-linear function.

### 4 Convolutional 2D Embeddings of Knowledge Graphs

In this work we propose a neural link prediction model where the interactions between input entities and relationships is modelled by fully-connected and convolutional layers. Our model’s main feature is convolution over 2D shaped embeddings. The architecture is summarised in Figure 1. Formally,

Figure 1: In the ConvE model the entity and relation embeddings are first reshaped and concatenated (steps 1, 2) and the resulting matrix is used as an input to a convolutional layer (step 3) the resulting feature map tensor is vectorised and projected in a  $k$ -dimensional space (step 4) and matched with all candidate object embeddings (step 5).



the scoring function is defined as follows:

$$\psi_r(\mathbf{e}_s, \mathbf{e}_o) \triangleq f(\text{vec}(f([\overline{\mathbf{e}}_s; \overline{\mathbf{r}}_r] * \omega))) \mathbf{W} \mathbf{e}_o, \quad (1)$$

where  $f$  denotes a non-linear function, and  $\overline{\mathbf{e}}_s$  and  $\overline{\mathbf{e}}_o$  denote a 2D reshaping of  $\mathbf{e}_s$  and  $\mathbf{e}_o$ , respectively: if  $\mathbf{e}_s, \mathbf{e}_o \in \mathbb{R}^k$ , then  $\overline{\mathbf{e}}_s, \overline{\mathbf{e}}_o \in \mathbb{R}^{k_w \times k_h}$ , where  $k = k_w k_h$ .

In the feed-forward pass, the model performs a row-vector look-up operation in two embedding matrices, one for entities, denoted  $\mathbf{E}^{|\mathcal{E}| \times k'}$  and one for relations, denoted  $\mathbf{R}^{|\mathcal{R}| \times k'}$ , where  $k$  and  $k'$  are respectively the entity and relation embedding dimensions, and  $|\mathcal{E}|$  and  $|\mathcal{R}|$  denote the number of entities and relations. The model then concatenates  $\overline{\mathbf{e}}_s$  and  $\overline{\mathbf{e}}_o$ , and uses it as an input for a 2D convolutional layer with filters  $\omega$ . Such a layer returns a feature map tensor  $\mathcal{T} \in \mathbb{R}^{c \times m \times n}$ , where  $c$  is the number of 2D filters, and  $m$  and  $n$  are the dimensions of the extracted feature maps. The tensor  $\mathcal{T}$  is then reshaped in a  $\text{vec}(\mathcal{T}) \in \mathbb{R}^{cmn}$  vector, which is then projected in a  $k$ -dimensional space by a linear transformation parametrised by the matrix  $\mathbf{W} \in \mathbb{R}^{cmn \times k}$  and matched with the object embedding  $\mathbf{e}_o$  via a dot product. The convolutional filters and the matrix  $\mathbf{W}$  are shared parameters, independent of the input entities  $s$  and  $o$  and the relationship  $r$ .

For training the model parameters, we apply a logistic sigmoid to the logits of the scores of  $(s, r, o)$ , and minimise the following binary cross-entropy loss:

$$\mathcal{L}(p, t) = -\frac{1}{N} \sum_i (t_i \cdot \log(p_i) + (1 - t_i) \cdot \log(1 - p_i)), \quad (2)$$

where  $p$  is the prediction and  $t$  the label.

We use rectified linear units as the non-linearity  $f$  for faster training [14], and batch normalisation after each layer to stabilise, regularise and increase rate of training convergence [9]. We regularise our model by using dropout [28] in several stages: We dropout the embeddings, feature maps after the convolution operation and hidden units after the fully connected layer. We use Adam as optimiser [12], and label smoothing to lessen overfitting due to saturation of output non-linearities at the labels [30].

#### 4.1 Fast Evaluation for Link Prediction Tasks

Unlike other link prediction models which take an entity pair and a relation as a triple  $(s, r, o)$ , and score it (1-1 approach), we take one  $(s, r)$  pair and score it against all entities  $o \in \mathcal{E}$  simultaneously (1-N approach). In our architecture convolution consumes about 75-90% of the total computation time of the model and thus it is important to minimise the number of convolution operations. If we take a naive 1-1 approach, then a training pass and an evaluation with a convolution model on FB15k – one of the dataset used in the experiments – takes 2.4 minutes and 3.34 hours, respectively, using a high-end GPU with a batch size of 128 and embedding size 128. Using a 1-N approach the respective numbers are 45 and 35 seconds – a considerable improvement of over 300x in terms of evaluation time. Usually, in a 1-1 approach, the batch size is increased to speed up evaluation [2], but this is not feasible for convolutional models since the GPU memory requirements quickly blows up for large

batch sizes when one uses convolution. Thus the 1-N evaluation is of practical importance, since our model would be computationally expensive to evaluate otherwise.

Do note that any 1-1 model can make use of 1-N evaluation. Thus this practical trick in speeding up the evaluation is applicable for all standard link prediction models which usually operate using a 1-1 approach.

## 5 Experiments

### 5.1 Knowledge Graph Datasets

For evaluating our proposed model, we use a selection of common knowledge graphs from the literature. WN18 [1] is a subset of WordNet which consists of 18 relations and 40,943 entities. Most of the 151,442 triples have hyponym and hypernym relations and thus WN18 is marked by relations arranged in a strictly hierarchical structure. FB15k [1] is a subset of Freebase which contains about 15k entities with 1,345 different relations. A large fraction of content in this knowledge graph deals with movies/actors/awards and sports/sport teams. YAGO3-10 [19] is a subset of YAGO3 which consists of entities which have a minimum of 10 relations each. It has 123,182 entities and 37 relations. Most of the triples deal with descriptive attributes of people (citizenship/gender/profession).

Countries is a benchmark dataset that is useful to evaluate a model’s ability to learn long-range dependencies between entities and relations. It consists of three sub-tasks which increase in difficulty in a step-wise fashion. It consists of three types of entities: countries (e.g. Germany, Italy), sub-regions (e.g. Western Europe, South America) and regions (e.g. Europe, Americas); and two relations: NeighborOf(country, country) and LocatedIn(country, sub-region / region). The task is to predict the region of a given country. It is evaluated in terms of area-under-the-curve for precision/recall (AUC-PR). The task has three levels of difficulty, S1, S2, S3, where each step increases the path-length or indirection which one needs to traverse the dataset graph in order to find a solution. In S1 all regions of test-case countries are removed from the training set (solution: Country  $\rightarrow$  sub-region  $\rightarrow$  region); in S2 one also removes the sub-regions (solution: Country  $\rightarrow$  neighbour  $\rightarrow$  region); in S3 one also removes the region of all the neighbors (solution: Country  $\rightarrow$  neighbour  $\rightarrow$  sub-region  $\rightarrow$  region).

We also introduce a new dataset, WN18RR,<sup>2</sup> which is an alteration of WN18 which aims to address the shortcoming of WN18: It was first noted by Toutanova and Chen [31] that the test sets of WN18 and FB15k contain mostly reversed triples that are present in the training set, for example the test set contains (*s*, hyponym, *o*) while the training set contains the reverse (*o*, hypernym, *s*). Toutanova and Chen [31] introduced FB15k-237, a subset of FB15k where reversing relations are removed, to create a dataset without this property. However, they did not explicitly demonstrate the severity of this problem, which might explain the fact that further research kept using these datasets for evaluation without addressing this issue.

In Section 5.2 we introduce a simple reversal rule which demonstrates the severity of this bias by achieving state-of-the-art results on both WN18 and FB15k. One might argue, that a good relational model should learn how to reverse relations in addition to more complex aspects of the dataset, but if this is our evaluation goal we should design more controlled datasets and experiments where it is clear what a relational model learns. By creating WN18RR, we seek to reclaim WN18 as a dataset which tests a models ability to model a general knowledge graph which cannot easily be completed using a single rule. As such, we do not recommend the usage of FB15k and WN18 in further research. Instead, we recommend the usage of FB15k-237, WN18RR, and YAGO3-10 that do not suffer from these issues.

### 5.2 Experimental Setup

We selected the hyperparameters of our ConvE model via grid search according to the mean reciprocal rank (MRR) on the validation set. Hyperparameter ranges for the grid search were the following – embedding dropout in {0.0, 0.1, 0.2}, feature map dropout in {0.0, 0.1, 0.2, 0.3}, projection layer dropout in {0.0, 0.1, 0.3, 0.5}, embedding size in {100, 200}, batch size in {64, 128, 256}, learning rate in {0.001, 0.003}, label smoothing in {0.0, 0.1, 0.2, 0.3}. Besides the grid search, we also tried modifications of the 2D convolution layer in our models: We tried replacing it with fully connected

---

<sup>2</sup>Code and dataset available at: <https://github.com/TimDettmers/ConvE>

layers, and 1D convolution, but these performed consistently worse and we abandoned them. We also experimented with different filter sizes and found that we only receive good results if the first convolutional layer uses small filters, that is 3x3 filters. We found that the following combination of parameters works well on WN18, YAGO3-10 and FB15k: embedding dropout 0.2, feature map dropout 0.2, projection layer dropout 0.3, embedding size 200, batch size 128, learning rate 0.001, label smoothing 0.1. For the Countries dataset, we increase embedding dropout to 0.3, hidden dropout to 0.5 and set label smoothing to 0.

We use early stopping using the mean reciprocal rank (WN18, FB15k, YAGO3-10) and AUC-PR (Countries) statistics on the validation set which we evaluate every three epochs. Unlike the other datasets, for Countries the results have a high variance, as such we average 10 runs and produce 95% confidence intervals.

**Baseline – Reversal model:** WN18 and FB15k has been noted to contain many reversible relations [31], for instance, a test triple (feline, *hyponym*, cat) can easily be mapped to a training triple (cat, *hyponym*, feline): knowing that hyponym is the inverse of hypernym allows you to easily predict the vast majority of test triples.

For such a reason, to test the redundancy of the datasets, we also investigated a simple rule-based baseline model, which we will refer to as the *reverse model*. This model looks for relationships which are the reverse of each other, such as hypernym and hyponym. We extract these relationships automatically from the training set: given two relation pairs  $r_1, r_2 \in \mathcal{R}$ , we check whether  $(s, r_1, o)$  implies  $(o, r_2, s)$ , and vice-versa. If the presence of  $(s, r_1, o)$  co-occurs with the presence of  $(o, r_2, s)$  at least 99% of the time, we say that  $r_1$  implies  $r_2$  ( $r_1 \implies r_2$ ) and we use this reversal rule for predicting test triples. We use the training set to check for reversal rules. At test time we check if the test triple has reversal matches outside the test set, if  $k$  matches are found we sample a permutation of the top  $k$  ranks for these matches; if no match is found we select a random rank for the test triple.

## 6 Results

Similarly to [34, 33, 24], we here focus on reporting results in a filtered setting, that is we rank triples only against scores for all possible entity combinations of unknown triples and we do not rank against combinations of existing, known triples. Our results on the standard benchmarks FB15k and WN18 are shown in Table 2, results on the datasets with reversing relations removed are shown in Table 3; results on YAGO3-10 and Countries are shown in Table 4.

Strikingly, the reverse model baseline achieves state-of-the-art on many different metrics on both, FB15k and WN18 datasets. However, it fails to pick up on reversible relations on YAGO3-10 and FB15k-237. On WN18 our reverse model achieves a good score due to self-reversing relationships like "similar to" which still exist in the dataset after using the procedure which is also used for the FB15k-237 dataset.

Our proposed model, ConvE, achieves state-of-the-art performance for all metrics on YAGO3-10, for some metrics on FB15k, and it does well on WN18. On Countries, it solves the S1 and S2 tasks, and does well on S3, scoring better than other models like DistMult and ComplEx

Table 2: Link prediction results on WN18 and FB15k

	WN18					FB15k				
	MR	MRR	@10	Hits @3	@1	MR	MRR	@10	Hits @3	@1
DistMult [34]	902	0.822	0.936	0.914	0.728	97	0.654	0.824	0.733	0.546
ComplEx [33]	–	0.941	0.947	0.936	0.936	–	0.692	0.840	0.759	0.599
Gaifman [24]	<b>352</b>	–	0.939	–	0.761	75	–	0.842	–	0.692
ANALOGY [18]	–	<b>0.942</b>	0.947	0.944	<b>0.939</b>	–	0.725	0.854	0.785	0.646
R-GCN [26]	–	0.814	0.964	0.929	0.697	–	0.696	0.842	0.760	0.601
ConvE	504	<b>0.942</b>	0.955	0.947	0.935	<b>64</b>	0.745	<b>0.873</b>	<b>0.801</b>	0.670
ReverseModel	602	0.857	<b>0.969</b>	<b>0.958</b>	0.757	1563	<b>0.759</b>	0.786	0.771	<b>0.743</b>

Table 3: Link prediction results on WN18RR and FB15k-237

	WN18RR					FB15k-237				
	MR	MRR	Hits			MR	MRR	Hits		
			@10	@3	@1			@10	@3	@1
DistMult [34]	<b>5110</b>	0.425	0.491	0.439	0.389	254	0.241	0.419	0.263	0.155
ComplEx [33]	5261	<b>0.444</b>	<b>0.507</b>	<b>0.458</b>	<b>0.411</b>	<b>248</b>	0.240	0.419	0.263	0.152
R-GCN [26]	–	–	–	–	–	–	0.248	0.417	0.258	0.153
ConvE	7323	0.342	0.411	0.360	0.306	330	<b>0.301</b>	<b>0.458</b>	<b>0.330</b>	<b>0.220</b>
ReverseModel	13417	0.360	0.360	0.360	0.360	7124	0.007	0.012	0.008	0.004

Table 4: Link prediction results on YAGO3-10 and Countries

	YAGO3-10					Countries		
	MR	MRR	Hits			AUC-PR		
			@10	@3	@1	S1	S2	S3
DistMult [34]	5926	0.337	0.540	0.379	0.237	<b>1.000±0.000</b>	0.721±0.122	0.516±0.070
ComplEx [33]	6351	0.355	0.547	0.399	0.258	0.965±0.021	0.571±0.104	0.430±0.072
ConvE	<b>2792</b>	<b>0.523</b>	<b>0.658</b>	<b>0.564</b>	<b>0.448</b>	<b>1.000±0.000</b>	<b>0.985±0.013</b>	<b>0.856 ±0.051</b>
ReverseModel	60251	0.015	0.022	0.017	0.010	–	–	–

For FB15k-237, we could not replicate the basic model results from Toutanova et al. [32], where the models in general have better performance than what we can achieve. Compared to Schlichtkrull et al. [26] our results for standard models are a bit better than theirs and on-a-par with their R-GCN model.

## 7 Analysis

### 7.1 Looking at Indegree and PageRank

Our main hypothesis for the good performance of our model on datasets like YAGO3-10 and FB15k-237 compared to WN18RR is that these dataset contain nodes with very high relation-specific indegree. For example the node "United States" (entity embedding) with edges "was born in" (relation embedding) has an indegree of over 10,000. Many of these 10,000 nodes will be very different from each other (actors, writers, politicians, business people) and our main hypothesis is that models that learn multiple layers of non-linear features like ours have an advantage over shallow models to capture all these constraints for such high indegree nodes.

However, for simpler datasets, like WN18 which mainly consists of hypernym/hyponym relations which often have an indegree of one (there is often only one generalisation for a given concept), we have nodes with small relation-specific indegree and thus a linear model might be sufficient, easier to optimise, and thus be able to find a better local minimum.


In this section we compare DistMult, a model that uses a simple tri-linear dot product, and our model, ConvE, that learn multiple layers of non-linear features to analyse the effect these high indegree nodes on performance.

We test our hypothesis in two ways. We remove triples which contain nodes which have a relation-specific indegree of greater than two for FB15k; and we remove triples which contain nodes which have a relation-specific indegree less than two for WN18. On these datasets we hypothesise that compared to DistMult, (1) our model will perform worse on FB15k (relatively more nodes with low indegree), and (2) our model will perform better on WN18 (relatively more nodes with high indegree). Indeed, we find that both hypotheses hold: For (1) on FB15k we have ConvE 0.586 Hits@10 vs DistMult 0.728 Hits@10; for (2) on WN18 we have ConvE 0.952 Hits@10 vs DistMult 0.938 Hits@10. This shows that our model indeed might have an advantage when modelling nodes with high indegree.

Table 5: Mean PageRank  $\times 10^{-3}$  of nodes in the test set vs reduction in error in terms of AUC-PR or Hits@10 of ConvE wrt. DistMult.

Dataset	PageRank	Error Reduction
WN18RR	0.104	0.91
WN18	0.125	1.28
FB15k	0.599	1.23
FB15RR	0.733	1.17
YAGO3-10	0.988	1.91
Countries S3	1.415	3.36
Countries S1	1.711	0
Countries S2	1.796	18.6

Figure 2: Visualisation of convolutional filters for (a) WN18 and (b) YAGO3-10.



To verify this hypothesis further we look at PageRank, which is a measure of centrality of a node. PageRank can also be seen as a measure of the recursive indegree of a node, that is, the PageRank value of a node is proportional to the indegree of this node, its neighbours indegree, its neighbours-neighbours indegree and so forth scaled relative to all other nodes in the network.

In line with our argument above, we expect that nodes with high PageRank are more difficult to model, since one entity embedding needs to capture numerous constraints with other entity embeddings, and additionally, many of its neighbours, which by definition of high PageRank have often high indegree, will need to capture numerous additional constraints and so forth.

To test this hypothesis, we calculate the PageRank for each dataset as a measure of centrality. We find that the most central nodes in WN18 have a PageRank value more than one order of magnitude smaller than the most central nodes in YAGO3-10 and Countries, and about 4 times smaller than the most central nodes in FB15k; see Figure 3 in the Appendix for more statistics. When we look at the mean PageRank of nodes contained in the test sets, we find that the difference of performance in terms of Hits@10 between DistMult and ConvE is roughly proportional to the mean test set PageRank, that is, the higher the mean PageRank of the test set nodes the better does ConvE compared to DistMult and vice-versa. See Table 5 for these statistics. The correlation between mean test set PageRank and relative error reduction of ConvE compared to DistMult is strong with  $r = 0.83$ . This gives additional evidence that our model has an advantage at modelling nodes with high (recursive) indegree.

To verify if this behaviour is caused by learning multiple layers of non-linear features, we remove the convolutional layer from our model and replace it with either two fully connected layers or with a 1D convolutional layer while keeping the rest of the architecture consistent. We then train all three models on FB15k and compare the scores. If multiple layers of non-linear features would be the decisive factor we would expect similar score for all three models. We find that the fully connect architecture performs much worse than 1D or 2D convolution, with Hits@10 0.466 (MLP) vs 0.821 (1D) vs 0.873 (2D), respectively.

Experimentally, this shows that the multiple layers of non-linear features are not the decisive factors, but rather the ability of a layer to extract useful features. This is similar to the performance of multi-layer perceptrons compared to convolutional networks in computer vision — both models learn non-linear features, but convolutional layers can extract more relevant features which lead to better generalisation. The differences between 1D and 2D convolution suggests that features learned from the interaction of two embeddings (2D) are more powerful than just modelling the embeddings with convolution (1D).

In conclusion, we believe that the increased performance of our model compared to a standard link predictor, DistMult, can be partially be explained due to our model’s ability to model nodes with high indegree with greater precision. We show that this behaviour might be due to the ability of 2D convolutional layers to learn features of interactions between entities and relation embeddings. This begs the question what kind of spatial structure 2D convolution is inducing in the embedding space.



## 7.2 Spatial Structure of 2D Embeddings

Here we test if the 2D embeddings contain spatial structure. We use Moran's I test which tests the null hypothesis that the global spatial autocorrelation is zero, or in other words, that the random variable does not have any global spatial structure [20]. We use the PySal package [25] to carry out the test.

We normalise each "image" to have mean zero and variance and apply the Moran's I test to up to 100 samples of entity and relation embeddings. We count the proportion of significant tests to get evidence if these embeddings have some structure in general. We find that Moran I tests on entity embeddings are significant near chance level, with less than 10% of tests being significant at the  $p = 0.05$  level. For relation embeddings we find weak evidence for spatial structure for YAGO3-10 where 44% of relations are significant. For visualisations of these significant relations see Figure 4 in the Appendix. These results suggest that 2D entity embeddings generally do not have global spatial structure and that only some of the relation embedding have global spatial structure.

However, visualisation of YAGO3-10 and WN18 convolution filters, as seen in Figure 2 suggest that these filters match contrasting "pixel" values, that is, the filters match patterns where one pixel or a group of pixels stands out relative to its neighbourhood. These pixel-feature filters give some hints why Moran's I test failed to pick up spatial patterns, since a spatial pattern in embeddings must be larger than a few contrasting pixels to be significantly different from noise. If the 2D convolutions would not pick up on any spatial features, we would expect that they lead to poor model performance so this means that these filters must work on pixel-level structures. This also explains why convolutions larger than 3x3 do not work well in the first convolutional layer since larger filters sum pixel-level features together with their surroundings thus diluting information in pixel-level structures.

## 8 Conclusion and Future Work

Here we introduced ConvE, a link prediction model that uses 2D convolution over embeddings and multiple layers of non-linear features to model knowledge graphs. This model uses few parameters and is computationally efficient due to a 1-N approach of link prediction and thus scales well with increasing knowledge graph sizes. Our model achieves state-of-the-art results on several existing knowledge graph datasets. However, building on previous work, we also show that a simple reverse model for relations can achieve state-of-the-art results on WN18 and FB15k thus questioning if models on this dataset actually learn general link prediction rather than learning this reversal rule. We introduce WN18RR to address this issue for WN18 and we recommend using FB15k-237 over FB15k for future research.

In our analysis we show that the performance of our model compared to a common link predictor, DistMult, can partially be explained by its ability to model nodes with high (recursive) indegree. Tests for spatial autocorrelation reveal that the entity and relation embeddings in general do not have any significant spatial structure, but that 2D convolutions on embeddings instead learn structures of contrasting pixels.

Our model is still shallow compared to convolutional architecture found in computer vision and future work might deal with convolutional models of increasing depth. Further work might also look at the interpretation of 2D convolution or how to enforce large-scale structure in embedding space and thus make convolutional filters learn feature extractors for large scale structures in embedding interactions.

### Acknowledgments

We would like to thank Johannes Welbl for his feedback and helpful discussions related to this work. This work was supported by a Marie Curie Career Integration Award, an Allen Distinguished Investigator Award, a Google Europe Scholarship for Students with Disabilities, and the H2020 project SUMMA.

### References

- [1] Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. Translating Embeddings for Modeling Multi-relational Data. In Christopher J. C. Burges et al.,

- editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 2787–2795, 2013.
- [2] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating Embeddings for Modeling Multi-relational Data. In *Advances in neural information processing systems*, pages 2787–2795, 2013.
  - [3] Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. A semantic matching energy function for learning with multi-relational data - application to word-sense disambiguation. *Machine Learning*, 94(2):233–259, 2014.
  - [4] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel P. Kuksa. Natural Language Processing (Almost) from Scratch. *Journal of Machine Learning Research*, 12:2493–2537, 2011.
  - [5] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In Lee et al. [17], pages 3837–3845.
  - [6] Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. Knowledge Vault: A Web-Scale Approach to Probabilistic Knowledge Fusion. In Sofus A. Macskassy et al., editors, *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, pages 601–610. ACM, 2014. ISBN 978-1-4503-2956-9.
  - [7] David K. Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. Convolutional Networks on Graphs for Learning Molecular Fingerprints. In Corinna Cortes et al., editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2224–2232, 2015.
  - [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
  - [9] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv preprint arXiv:1502.03167*, 2015.
  - [10] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. A Convolutional Neural Network for Modelling Sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 655–665. The Association for Computer Linguistics, 2014. ISBN 978-1-937284-72-5.
  - [11] Yoon Kim. Convolutional Neural Networks for Sentence Classification. In Alessandro Moschitti et al., editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751. ACL, 2014. ISBN 978-1-937284-96-1.
  - [12] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
  - [13] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*, 2016.
  - [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
  - [15] Denis Krompaß, Stephan Baier, and Volker Tresp. Type-Constrained Representation Learning in Knowledge Graphs. In Marcelo Arenas et al., editors, *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part I*, volume 9366 of *Lecture Notes in Computer Science*, pages 640–655. Springer, 2015. ISBN 978-3-319-25006-9.
  - [16] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.

- [17] Daniel D. Lee et al., editors. *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 2016.
- [18] H. Liu, Y. Wu, and Y. Yang. Analogical Inference for Multi-Relational Embeddings. *ArXiv e-prints*, May 2017.
- [19] Farzaneh Mahdisoltani, Joanna Biega, and Fabian M. Suchanek. YAGO3: A Knowledge Base from Multilingual Wikipedias. In *CIDR 2015, Seventh Biennial Conference on Innovative Data Systems Research, Asilomar, CA, USA, January 4-7, 2015, Online Proceedings*. www.cidrdb.org, 2015.
- [20] Patrick AP Moran. Notes on Continuous Stochastic Phenomena. *Biometrika*, 37(1/2):17–23, 1950.
- [21] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A Three-Way Model for Collective Learning on Multi-Relational Data. In Lise Getoor et al., editors, *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 809–816. Omnipress, 2011.
- [22] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2016.
- [23] Maximilian Nickel, Lorenzo Rosasco, and Tomaso A. Poggio. Holographic Embeddings of Knowledge Graphs. In Dale Schuurmans et al., editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 1955–1961. AAAI Press, 2016. ISBN 978-1-57735-760-5.
- [24] Mathias Niepert. Discriminative Gaifman Models. In Lee et al. [17], pages 3405–3413.
- [25] Sergio J Rey and Luc Anselin. PySAL: A Python Library of Spatial Analytical Methods. *Handbook of applied spatial analysis*, pages 175–193, 2010.
- [26] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks. *arXiv preprint arXiv:1703.06103*, 2017.
- [27] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. Learning Semantic Representations Using Convolutional Neural Networks for Web Search. In Chin-Wan Chung et al., editors, *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014, Companion Volume*, pages 373–374. ACM, 2014. ISBN 978-1-4503-2745-9.
- [28] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [29] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
- [30] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [31] Kristina Toutanova and Danqi Chen. 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, 2015.
- [32] Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoifung Poon, Pallavi Choudhury, and Michael Gamon. Representing Text for Joint Embedding of Text and Knowledge Bases. In *EMNLP*, volume 15, pages 1499–1509, 2015.
- [33] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex Embeddings for Simple Link Prediction. In Maria-Florina Balcan et al., editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org, 2016.

- [34] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *International Conference on Learning Representations (ICLR)*, 2015.
- [35] Wen-tau Yih, Kristina Toutanova, John C. Platt, and Christopher Meek. Learning Discriminative Projections for Text Similarity Measures. In Sharon Goldwater et al., editors, *Proceedings of the Fifteenth Conference on Computational Natural Language Learning, CoNLL 2011, Portland, Oregon, USA, June 23-24, 2011*, pages 247–256. ACL, 2011. ISBN 978-1-932432-92-3.

## 9 Appendix

Figure 3: PageRank and other statistics for nodes in the Countries S3, YAGO3-10, FB15k and WN18 datasets

0, americas	degree: 47,	coverage: 0.045,	aggregated coverage: 0.045,	PageRank: 0.037
1, oceania	degree: 30,	coverage: 0.029,	aggregated coverage: 0.075,	PageRank: 0.027
2, asia	degree: 40,	coverage: 0.039,	aggregated coverage: 0.11,	PageRank: 0.023
3, caribbean	degree: 28,	coverage: 0.027,	aggregated coverage: 0.14,	PageRank: 0.023
4, europe	degree: 35,	coverage: 0.034,	aggregated coverage: 0.17,	PageRank: 0.022
5, africa	degree: 36,	coverage: 0.035,	aggregated coverage: 0.21,	PageRank: 0.021
6, northern_europe	degree: 16,	coverage: 0.015,	aggregated coverage: 0.22,	PageRank: 0.011
7, polynesia	degree: 11,	coverage: 0.011,	aggregated coverage: 0.24,	PageRank: 0.01
8, china	degree: 17,	coverage: 0.016,	aggregated coverage: 0.25,	PageRank: 0.0097
9, eastern_africa	degree: 15,	coverage: 0.015,	aggregated coverage: 0.27,	PageRank: 0.0095
10, russia	degree: 15,	coverage: 0.015,	aggregated coverage: 0.28,	PageRank: 0.0087

0, male	degree: 61604,	coverage: 0.057,	aggregated coverage: 0.057,	PageRank: 0.032
1, United_States	degree: 12263,	coverage: 0.011,	aggregated coverage: 0.068,	PageRank: 0.0099
2, female	degree: 5167,	coverage: 0.0047,	aggregated coverage: 0.073,	PageRank: 0.0039
3, United_Kingdom	degree: 1854,	coverage: 0.0017,	aggregated coverage: 0.074,	PageRank: 0.0013
4, wordnet_guitar_103467517	degree: 1673,	coverage: 0.0015,	aggregated coverage: 0.076,	PageRank: 0.0011
5, France	degree: 1581,	coverage: 0.0015,	aggregated coverage: 0.077,	PageRank: 0.0011
6, London	degree: 1643,	coverage: 0.0015,	aggregated coverage: 0.079,	PageRank: 0.001
7, New_York_City	degree: 1393,	coverage: 0.0013,	aggregated coverage: 0.08,	PageRank: 0.00093
8, India	degree: 1042,	coverage: 0.00096,	aggregated coverage: 0.081,	PageRank: 0.00092
9, Los_Angeles	degree: 1350,	coverage: 0.0012,	aggregated coverage: 0.082,	PageRank: 0.0009
10, Poland	degree: 834,	coverage: 0.00077,	aggregated coverage: 0.083,	PageRank: 0.00083

0, United States of America	degree: 7071,	coverage: 0.012,	aggregated coverage: 0.012,	PageRank: 0.012
1, /m/08mbj5d	degree: 4377,	coverage: 0.0074,	aggregated coverage: 0.019,	PageRank: 0.0093
2, /m/04ztj	degree: 3682,	coverage: 0.0062,	aggregated coverage: 0.026,	PageRank: 0.0067
3, /m/05zppz	degree: 3639,	coverage: 0.0061,	aggregated coverage: 0.032,	PageRank: 0.0065
4, /m/09nqf	degree: 3291,	coverage: 0.0056,	aggregated coverage: 0.037,	PageRank: 0.0056
5, /m/08mbj32	degree: 2868,	coverage: 0.0048,	aggregated coverage: 0.042,	PageRank: 0.0055
6, English language	degree: 3206,	coverage: 0.0054,	aggregated coverage: 0.048,	PageRank: 0.0049
7, /m/02hrh1q	degree: 2771,	coverage: 0.0047,	aggregated coverage: 0.052,	PageRank: 0.0046
8, /m/02nzb8	degree: 763,	coverage: 0.0013,	aggregated coverage: 0.053,	PageRank: 0.0033
9, United Kingdom	degree: 1770,	coverage: 0.003,	aggregated coverage: 0.056,	PageRank: 0.0033
10, /m/02sdk9v	degree: 763,	coverage: 0.0013,	aggregated coverage: 0.058,	PageRank: 0.0033

0, __united_kingdom_of_great_britain_and_northern_ireland_NN_1	degree: 498,	coverage: 0.0033,	aggregated coverage: 0.0033,	PageRank: 0.003
1, __soul_NN_2	degree: 411,	coverage: 0.0027,	aggregated coverage: 0.006,	PageRank: 0.0026
2, __urban_center_NN_1	degree: 520,	coverage: 0.0034,	aggregated coverage: 0.0094,	PageRank: 0.0024
3, __war_machine_NN_1	degree: 366,	coverage: 0.0024,	aggregated coverage: 0.012,	PageRank: 0.002
4, __bird_genus_NN_1	degree: 295,	coverage: 0.0019,	aggregated coverage: 0.014,	PageRank: 0.0016
5, __asterid_dicot_genus_NN_1	degree: 306,	coverage: 0.002,	aggregated coverage: 0.016,	PageRank: 0.0016
6, __herbaceous_plant_NN_1	degree: 251,	coverage: 0.0017,	aggregated coverage: 0.017,	PageRank: 0.0016
7, __modify_VB_3	degree: 275,	coverage: 0.0018,	aggregated coverage: 0.019,	PageRank: 0.0016
8, __law_NN_1	degree: 299,	coverage: 0.002,	aggregated coverage: 0.021,	PageRank: 0.0016
9, __mammal_genus_NN_1	degree: 265,	coverage: 0.0017,	aggregated coverage: 0.023,	PageRank: 0.0015
10, __trade_name_NN_1	degree: 245,	coverage: 0.0016,	aggregated coverage: 0.025,	PageRank: 0.0014

Figure 4: Visualisations of normalised YAGO3-10 relations which were significant under Moran's I test for spatial auto-correlation.

