

---

# Relational Pattern Benchmarking on the Knowledge Graph Link Prediction Task

---

Afshin Sadeghi\*<sup>1,2</sup>, Hirra Abdul Malik<sup>1</sup>, Diego Collarana<sup>2,3</sup>, and Jens Lehmann<sup>1,2</sup>

<sup>1</sup>SDA Research Group, University of Bonn, Germany

<sup>2</sup>Fraunhofer IAIS, Dresden, Germany

<sup>3</sup>Universidad Privada Boliviana, Bolivia

{s6hiabdu,lehmann}.uni-bonn.de,

{afshin.sadeghi,diego.collarana.vargas,jens.lehmann}@iais.fraunhofer.de

## Abstract

Knowledge graphs (KGs) encode facts about the world in a graph data structure where entities, represented as nodes, connect via relationships, acting as edges. KGs are widely used in Machine Learning, e.g., to solve Natural Language Processing based tasks. Despite all the advancements in KGs, they plummet when it comes to completeness. Link Prediction based on KG embeddings targets the sparsity and incompleteness of KGs. Available datasets for Link Prediction do not consider different graph patterns, making it difficult to measure the performance of link prediction models on different KG settings. This paper presents a diverse set of pragmatic datasets to facilitate flexible and problem-tailored Link Prediction and Knowledge Graph Embeddings research. We define graph relational patterns, from being entirely inductive in one set to being transductive in the other. For each dataset, we provide uniform evaluation metrics. We analyze the models over our datasets to compare the model’s capabilities on a specific dataset type. Our analysis of datasets over state-of-the-art models provides a better insight into the suitable parameters for each situation, optimizing the KG-embedding-based systems.

**Keywords:** Machine Learning, Knowledge Graphs Embedding, Link Prediction, Benchmarking, Dataset, Relational Pattern, Inductive, Transductive

## 1 Introduction

Knowledge graphs constitute a significant part of NLP since the 70s. However, after the announcement of big hubs such as Google, Facebook, and Microsoft in the 90s, the growth of research in this particular field became evident as states Hogan et al. (2021). The different elements of a knowledge graph are represented in the form of triplets  $(h, r, t)$  where  $(h, t)$  represents entities a.k.a. ‘nodes’, whereas  $r$  tells the relationship between them, which is also known as ‘edge’. According to Yang and Mitchell (2017), NLP has gotten a new scope after the advancements in the field of knowledge graphs, easing the communication with machines. Various applications such as *information retrieval* Xiong et al. (2017), *question answering* Hao et al. (2017) and *recommender systems* Zhang et al. (2016) use knowledge graphs to improve their performance. Some of the most used knowledge graphs include *DBpedia* Lehmann et al. (2014), *Yago* Suchanek et al. (2007), *Freebase* Bollacker et al. (2008) and *WordNet* Miller (1995).

Despite being in demand, KGs still face many issues, such as data incompleteness. To tackle the issue, link prediction models observe the patterns in knowledge graphs based on how facts are

---

\*The two preceding authors are co-first.

connected together. According to Wang et al. (2019), the goal of the link prediction task is to map the entities/relations to low dimensional vectors capturing the structure of the knowledge graph, which helps predict the likelihood score of the triple. Despite advancements in benchmarks, a significant chunk is still unexplored. In this paper, we enhanced the work of Sarkar et al. (2020) by building various datasets on the principles of known and valuable facts using the Freebase and Wordnet datasets, categorizing them into different patterns for benchmarking. The relation of each triplet is observed and then grouped into categories. The categories include different patterns, e.g., symmetry has the ‘same’ relation between two entities such as ‘friends’, whereas inverse has two ‘different and directed’ relations between the entities such as ‘father and son’. We make categories of datasets and observe the link prediction (LP) models over them.

With the goal to set up a benchmark that separates the task of testing KG embedding models from the models, we extend the MLwin-Hobbit platform for benchmarking various trained methods, which are implemented in different environments (e.g. PyTorch and Java). This extension is crucial for a reproducible evaluation and fair comparison of methods.

We train state-of-the-art KG embedding models from scratch with our RAW datasets, providing experimental results that give exciting research directions. We consider the most popular and unified evaluation metrics along with the AUC-PR test in all combinations. Our experiments suggest that the link predicting models are scalable to large-scale datasets and graphs. These results indicate fruitful guidance for future research in KG Link Prediction and KG Embeddings.

### 1.1 Our Contribution

We propose several datasets<sup>2</sup> by classifying triplets into their respective classes according to their patterns, keeping in mind the properties from both inductive and transductive types. Therefore, we extract four categories from each class: Fully Inductive, Fully Transductive, CountBased Inductive, and either Head or Tail Inductive. Each category is further divided into patterns of Symmetry, Anti-symmetry, Inverse and Inductive, making a total of **sixteen** datasets. The datasets are also designed on the basis of unification to benchmark them onto different link predicting models. A significant setback in the benchmarking of knowledge graph models was observed. Therefore the work done by Hu et al. (2021) has been extended keeping in mind our set of data. A fair comparison is generated to help choose the best model and dataset combination for NLP Research. The previous research on benchmarking datasets was too general, we provide the specific approach to the type of datasets. We explore the characteristics of the datasets that can be potential performance boosters. To sum up our contribution, we created a reproducible evaluation environment that is user-friendly for all. We designed our benchmark datasets keeping the ease of user in mind.




Figure 1: Example of triplet categorization.

## 2 Related Work

**Dataset Building** We build our datasets keeping in mind the dataset building strategies from OGB (Hu et al. (2021)) as well as the CODEX (Safavi and Koutra (2020)) and TU-Dataset (Morris et al. (2020)). CODEX is gathered from thirteen domains ranging from medicine to music. This dataset is build on the principles of snowball sampling to extract data, while in our study, we searched for

<sup>2</sup>All datasets, scripts and extended results are available: [https://github.com/mlwin-de/relational\\_pattern\\_benchmarking](https://github.com/mlwin-de/relational_pattern_benchmarking)

specific relational patterns to enable on pattern-specific evaluations. TU-Dataset, a unified set of over 120 datasets from several domains, targets graph classification and regression tasks, while in our study, we focus on the link prediction task. We use the characteristics visualization technique for the datasets and the required properties for the characteristics analysis, of the aforementioned studies, as a baseline to define our dataset.

**Benchmarking** Benchmarking datasets help to compare and evaluate the LP models, parameters, and procedures as well as the statistics of the different datasets that are evaluated. The benchmarked dataset as described by Rossi et al. (2021) is useful for two basic analyses, efficiency and effectiveness analysis. The dataset of Hu et al. (2021) is kept as our ultimate guidance and standard to support the benchmarking task. CODEX dataset was benchmarked by Safavi and Koutra (2020) with unified evaluation strategies and empirical analysis. Relational Patterns of Inversion, Symmetry, and Composition were studied and subsets were prepared accordingly. We use the same strategy of making sub-datasets and then benchmarking on a number of models. Evaluation technique of using MRR and Hit Ratios was considered but along with the introduction of **AUC-PR**. Using GNN and Graph Kernel methods as used by Morris et al. (2020) gave us a new direction to use Teru et al. (2020) GraIL for our dataset benchmarking. In the study of relation patterns, we include the most frequent patterns. In the related works, they are studies that consider experiments on more complex logical rules, such as Meilicke et al. (2018) that evaluate on Inverse Equivalence and Subsumption rules, and in this direction, Yang et al. (2017) evaluates the performance of the knowledge base inference methods on a dataset of grid paths of different lengths.

In order to perform a link prediction benchmarking based on FAIR principles, we integrated the models in BenchEmbedd (Sadeghi et al. (2021b)), a platform that aims at benchmarking big linked data. Such benchmark experiments are frozen into docker containers, which can be accessed, reproduced, and reused easily with little prior knowledge of the test platform. The system allows researchers to make and test systems without having to worry about standardized hardware. Our trained models and datasets will be publicly available as a BenchEmbedd platform<sup>3</sup> for anyone to use and benchmark systems.

### 3 Dataset Building

Relation Extraction, the sub-field of information extraction, is one of the core techniques that support ML research. It organizes the structural information into groups according to the need (Wang et al. (2021)). We based our study on the two following datasets and extracted subsequent sub-datasets in the form of the stated patterns.

- **FB15K** A freebase dataset with a total of 592,213 triplets with 14,951 entities and 1,345 relationships. This factual dataset contains 483,142 Train triplets, 59,071 Test triplets and 50,000 Valid triplets. It dataset contains many entities from the wiki-link data.
- **WN18** A dataset extracted from Wordnet version 3 with a total of 141,442 triplets with 40,943 entities and 18 relationships. The dataset contains 141,442 Train triplet, 5,000 Test triplets and 5,000 Valid triplets. The dataset supports text analysis and provides with dictionary/thesaurus. Lexical relationship between synsets are stated by this dataset.

#### 3.1 Relational Datasets

Pattern extraction is the core task in building datasets for machine learning research. Pattern type suggests the type of link prediction model that works best for the given dataset. Patterns, that are also expressed as rules, each have different suitability to the embedding models. Link Predictors learn the specific pattern of the datasets and then match rule-based patterns to provide reasoning. The patterns we considered to build our datasets are stated below:

##### 3.1.1 Symmetry

This relational pattern is a sub-category of **Equivalence** pattern. Therefore, it is a binary relation that works in both directions. The relation can also be stated by the **Equal to** property, for instance, if

---

<sup>3</sup><https://github.com/mlwin-de/BenchEmbedd>

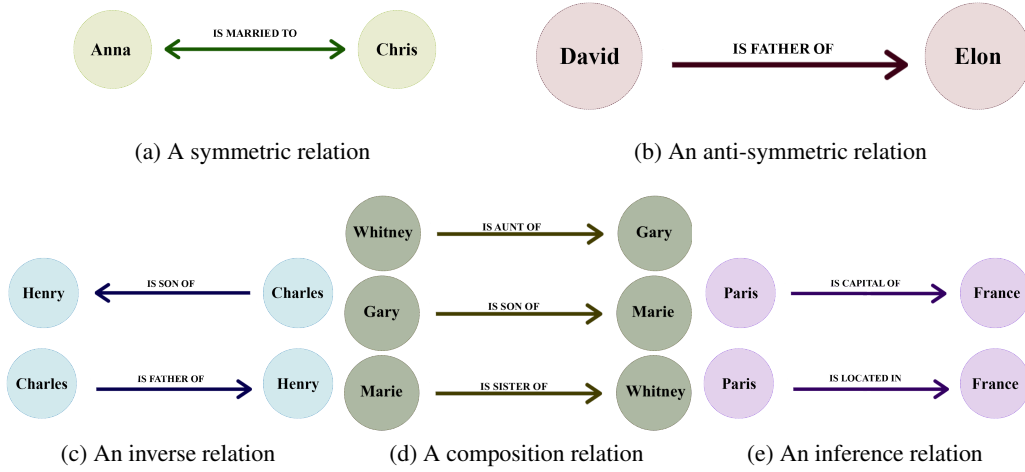


Figure 2: We define five different relational patterns to create our datasets and benchmark.

$a=b$  then  $b=a$  A relation is symmetric if:

$$\forall a, b \in X (aRb \Leftrightarrow bRa)$$

$$r(a, b) \Rightarrow r(b, a)$$

If  $R^T$  represents the converse of  $R$ , then  $R$  is symmetric if and only if  $R = R^T$ . Marriage, Friendship, and Partners are a few examples of symmetric relations.

### 3.1.2 Anti-Symmetry

The rule of Anti-Symmetry is opposite to Symmetry. It is a directed rule that states if a relation  $R$  binds  $A$  to  $B$ , the same can not work in the opposite direction, binding  $B$  to  $A$ . The rule is written as

$$r(a, b) \Rightarrow \neg r(b, a)$$

Relations such as Owner (to Tenant), Parent (to Child), and Singer (to Song) are Anti-Symmetric. Figure 2b shows a few examples of this rule.

### 3.1.3 Inverse

It is a binary relation stating two opposite relations for a set of entities. It is possible to assume a unique inverse relation for every relation. Inverse pattern between two set of two triples occurs when they have the opposite relation directions and have the same entities.

$$r_2(x, y) \Rightarrow r_1(y, x)$$

Parent-Child and Teacher-Student are examples of the inverse relations. Figure 2c shows how inverse relation is represented.

### 3.1.4 Composition

This binary relation which is also termed as **relation multiplication** is basically a **compound** relation which states the relation that can not exist without the existence of another relation. For Example, for the relation *Aunt*, the relation of *Sister* and *Son/Daughter* must exist in order to prove someone as an Aunt of somebody. In mathematical terms, relation  $r_1$  is composed of relation  $r_2$  and relation  $r_3$  if:

$$\forall x, y, z : r_2(x, y) \wedge r_3(y, z) \Rightarrow r_1(x, z)$$

Figure 2d states that Paris is the Capital of France, According to the composition property, Paris must be in France to be its capital.

Table 1: Inductive Setting Datasets

Inductive		Train	Test	Valid	Total Entities	Total Relations	% of Original Dataset
FB15K	Symmetry	2069	6	12	1062	3	0.3525
	Anti-Symmetry	9795	171	136	8539	529	1.7058
	Inverse	19590	373	298	8614	626	3.4213
	Inference	2634	98	43	1413	102	0.4670
WN18	Symmetry	31094	4	12	16991	3	20.5425
	Anti-Symmetry	8951	95	97	13129	18	6.0372
	Inverse	17902	190	196	13130	18	12.0759
	Inference	18000	944	962	13213	14	13.144

### 3.1.5 Inference

Inference relation pattern is one of the **logical** rules we formed datasets of. The rule states that we can deduce a relationship between two entities from the knowledge of another relationship between the two of them:

$$\forall h, t : (h, r_1, t) \Rightarrow (h, r_2, t)$$

where  $h$  and  $t$  are entities and  $r_1$  and  $r_2$  are relations between them. Figure 2e gives an example of inference relation which states Paris is the capital of France and thus, according to  $r_1$ , Paris must be in France as well. Therefore,  $(Paris, isCapitalof, France) \Rightarrow (Paris, isLocatedin, France)$

## 3.2 Evaluation Dataset Setting

There exist two different methods for the dataset division into train and test/validation subsets, where the composition of entities of each setting defines the evaluation setting:

**Inductive Setting** In an inductive setting, the entities during training are not found in the test dataset. The part of entities are kept missing and their relations are made to be found by the LP models. The number of disjoint entities varies in an inductively set dataset, fully disjoint sets are fully inductive and thus difficult for the models to predict.

**Transductive Setting** A dataset is divided to train and test/validation in a transductive setting when the occurrence of entities is ensured to be in the training procedure if it appears either in a test or valid sets. Transductive set datasets are best for Entity Specific Embedding. All entities in the training set are present in the test set and thus a model has trained embeddings for them specifically.

## 4 A New Set of Pattern Specific Datasets

### 4.1 The Standard Patterns

We extracted a subset of data from the standard FB15K and WN18 with Symmetry, Inverse, Anti-symmetry, and Inference patterns. Then, from each subset, we extracted an entirely inductive and a transductive dataset and two more customized datasets. In the first set, the percentage of inductive and transductive triples is fixed, and in the second set, each triplet has one Inductive entity with the other entity being transductive. The description of these settings is in the following.

**Inductive** We built<sup>4</sup> four datasets with the inductive setting where the entities of test and train datasets are entirely disjoint. We took these fully disjoint sets from both FB15K and WN18 and then subcategorized them into relational datasets, making a set of eight datasets. Table 1 states the statistics of our datasets.

**Transductive** Transductive Setting has common entities in train and test datasets. Therefore, the entities are already seen by the model, making prediction much easier for them. Table 2 states the statistics of the set of eight datasets from the transductive type.

<sup>4</sup>The script to extract data based on the relational patterns is available in [https://github.com/mlwin-de/relational\\_pattern\\_benchmarking/blob/master/utils/relationalpatterns.py](https://github.com/mlwin-de/relational_pattern_benchmarking/blob/master/utils/relationalpatterns.py)

Table 2: Transductive Setting Datasets

Transductive		Train	Test	Valid	Total Entities	Total Relations	% of Original Dataset
FB15K	Symmetry	2069	251	224	1044	3	0.42957
	Anti-Symmetry	9795	2083	1477	8005	529	2.2551
	Inverse	19590	4220	2983	8005	626	4.52425
	Inference	2634	112	74	1337	102	0.47618
WN18	Symmetry	31094	100	1074	16975	3	21.9674
	Anti-Symmetry	8951	53	54	12748	18	5.9811
	Inverse	17902	106	110	12748	18	11.9636
	Inference	18000	418	468	11549	14	12.4708

Table 3: Head-Tail Inductive Setting Datasets

Head-Tail Inductive		Train	Test	Valid	Total Entities	Total Relations	% of Original Dataset
FB15K	Symmetry	2069	12	6	1053	3	0.35071
	Anti-Symmetry	9795	1186	897	9605	529	1.84325
	Inverse	19590	2409	1825	9643	626	3.667098
	Inference	2634	184	152	1423	102	0.47314
WN18	Symmetry	31094	29	36	17012	3	20.5748
	Anti-Symmetry	8951	189	186	13121	18	20.7676
	Inverse	17902	378	372	13121	18	21.0271
	Inference	18000	1978	1940	13362	14	13.7293

**Head-Tail Ratio Inductive** We built a set of datasets by keeping either the head or the tail of each triplet in the train hidden from the test dataset. By doing so, we gain a semi-inductive dataset with each triplet unseen. Table 3 reports the statistics of these datasets.

**Percentage-wise building** In our study, we generated half of the test triples with inductive settings and half with the transductive setting. We apply this percentage base data generation on each category. Table 4 describes the statistics of the 50% datasets.

For our benchmark, we take standard evaluation metrics of Hit Ratios (at 1, 3, 10), Mean Reciprocal Rank, Area Under the Curve, and AUC-PR. Yousef et al. (2008) suggests, for a perfect AUC-PR score, an equal number of negative triplets are needed along with positive triplets. Therefore, as described by Teru et al. (2020), in the test set, the same number of negative samples are created by corrupting the copy of each triplet by either replacing the head or the tail with any random entity. We used the same procedure to incorporate each model with the AUC-PR score in a unified way. We considered **DistMult** Yang et al. (2015), **RotatE** Sun et al. (2019), **TransE** Bordes et al. (2013), **GraIL** Teru et al. (2020), **MDE** Sadeghi et al. (2020) and **CompGCN** Vashishth et al. (2020) for our analysis.

**Experiment Setup** Our system is implemented in Python, with Adadelta Zeiler (2012) as the optimizer. All Transductive bias models are set with learning rate **0.0001** with GraIL and CompGCN AT **0.01** in-order to practise uniformity. Alpha ( $\alpha$ ) is kept between  $[0.5, 1]$ . The models are set with the Dimensions =  $[GraIL = 1000, CompGCN = 100 \text{ and all other models} = 500]$ . The epochs are set uniform for TransE, Distmult and RotatE = **6000** whereas MDE is given a higher number of **150,000** and GraIL and CompGCN are run are **100** and **500** epochs respectively. For TransE, Distmult and RotatE, and MDE, we used a fixed number of negative samples, 50 in all the experiments. To regulate loss function, we estimated the score for 5 runs and took the average. All the experiments

Table 4: 50% Inductive Setting Datasets

Percentage Based Inductive		Train	Test	Valid	Total Entities	Total Relations	% of Original Dataset
FB15K	Symmetry	2069	12	24	1062	3	0.35544
	Anti-Symmetry	9795	342	272	8539	529	1.7576
	Inverse	19590	746	596	8614	626	3.5345
	Inference	2634	196	86	1413	102	0.49239
WN18	Symmetry	31094	8	24	16991	3	20.5530
	Anti-Symmetry	8951	148	151	13129	18	6.1079
	Inverse	17902	296	306	13130	18	12.2185
	Inference	18000	1362	1430	13213	14	13.7293

Table 5: Hit@10 and MRR results of Link Predictors on datasets extracted from FB15K.

Type of dataset	Dataset	FB15K													
		DistMult		TransE		RotatE		MDE		GraIL		CompGCN		QuatE	
		Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR
Inductive	Symm	0.0000	0.0015	0.0000	0.0012	0.0000	0.0009	0.0000	0.0230	0.0000	0.0366	0.0000	0.0187	0.0000	0.0040
	Anti-Sym	0.0000	0.0003	0.0000	0.0002	0.0000	0.0002	0.0000	0.0002	0.0000	0.2789	0.0000	0.0008	0.1986	0.0810
	Inverse	0.0000	0.0002	0.0000	0.0001	0.0000	0.0002	0.0000	0.0003	0.0000	0.0234	0.0000	0.0009	0.2252	0.0832
Transductive	Inference	0.1224	0.0544	0.1633	0.1640	0.1633	0.1639	0.0000	0.0007	0.0000	0.0260	0.0000	0.0039	0.2040	0.0934
	Symm	0.8625	0.8684	0.8944	0.3874	0.8645	0.8517	0.8486	0.7981	0.8267	0.8627	0.8685	0.8269	0.8804	0.8314
	Anti-Sym	0.0149	0.0084	0.1382	0.0484	0.0113	0.0066	0.0019	0.0010	0.2299	0.1866	0.1676	0.0950	0.3214	0.1695
Head/Tail Ratio	Inverse	0.0509	0.0352	0.1459	0.0540	0.0475	0.0346	0.0008	0.0009	0.2704	0.2062	0.1793	0.1030	0.2804	0.1501
	Inference	0.0358	0.0327	0.1161	0.0518	0.0357	0.0279	0.0267	0.0241	0.1298	0.1454	0.0536	0.0441	0.2009	0.1002
	Symm	0.0000	0.0021	0.0000	0.0023	0.0000	0.0037	0.0000	0.0083	0.0000	0.0270	0.0833	0.0903	0.0000	0.0066
Percentage Based (50%)	Anti-Sym	0.0038	0.0020	0.0868	0.0302	0.0055	0.0029	0.0008	0.0004	0.0046	0.0297	0.1749	0.0953	0.3473	0.1702
	Inverse	0.0033	0.0020	0.0849	0.0278	0.0061	0.0031	0.0000	0.0003	0.0017	0.0255	0.1814	0.0945	0.2929	0.1510
	Inference	0.0000	0.0016	0.1415	0.0616	0.0163	0.0066	0.0000	0.0015	0.0000	0.0222	0.0870	0.0630	0.3016	0.1687
Percentage Based (50%)	Symm	0.5000	0.4556	0.0833	0.0147	0.5000	0.4796	0.3758	0.1777	0.5000	0.5600	0.4167	0.3098	0.5000	0.4605
	Anti-Sym	0.0132	0.0101	0.0761	0.0246	0.0073	0.0038	0.0015	0.0008	0.1069	0.1038	0.0863	0.0450	0.3377	0.1457
	Inverse	0.0022	0.0138	0.0697	0.0245	0.0135	0.0089	0.0006	0.0005	0.1295	0.1101	0.0737	0.0438	0.2627	0.1212
Percentage Based (50%)	Inference	0.0818	0.0463	0.1224	0.1065	0.1428	0.1212	0.0204	0.0086	0.0714	0.0898	0.0563	0.0568	0.1658	0.0974

Table 6: Hit@10 and MRR results of Link Predictors on datasets extracted from WN18.

Type of dataset	Dataset	WN18													
		DistMult		TransE		RotatE		MDE		GraIL		CompGCN		QuatE	
		Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR	Hit@10	MRR
Inductive	Symm	0.0000	0.0002	0.0000	0.0001	0.0000	0.0001	0.0000	0.0001	0.0000	0.0200	0.0000	0.0019	0.0000	0.0000
	Anti-Sym	0.0000	0.0002	0.0000	0.0001	0.0000	0.0001	0.0000	0.0001	0.0000	0.0307	0.0000	0.0006	0.0000	0.0017
	Inverse	0.0000	0.0002	0.0000	0.0001	0.0000	0.0002	0.0000	0.0001	0.0000	0.0332	0.0000	0.0021	0.0000	0.0010
Transductive	Inference	0.0000	0.0002	0.0000	0.0001	0.0000	0.0001	0.0000	0.0001	0.0000	0.0399	0.0000	0.0015	0.0000	0.0002
	Symm	0.9722	0.9051	0.9713	0.4557	0.9732	0.9635	0.0014	0.0011	0.9755	0.9755	0.9722	0.9642	0.9727	0.9599
	Anti-Sym	0.0094	0.0022	0.0283	0.0113	0.0000	0.0007	0.0094	0.0043	0.0189	0.0338	0.0189	0.0060	0.0377	0.0153
Head/Tail Ratio	Inverse	0.0000	0.0012	0.0613	0.0155	0.0000	0.0007	0.0010	0.0010	0.0189	0.0455	0.0189	0.0038	0.0377	0.0156
	Inference	0.0335	0.0202	0.0467	0.0118	0.0287	0.0190	0.0019	0.0014	0.0287	0.0557	0.0311	0.0167	0.0586	0.0323
	Symm	0.0000	0.0004	0.0000	0.0001	0.0000	0.0002	0.0000	0.0001	0.0000	0.0284	0.0345	0.0096	0.0172	0.0043
Percentage Based (50%)	Anti-Sym	0.0053	0.0020	0.0291	0.0104	0.0000	0.0007	0.0000	0.0002	0.0000	0.0274	0.0396	0.0268	0.0534	0.0253
	Inverse	0.0040	0.0015	0.0185	0.0063	0.0000	0.0006	0.0000	0.0001	0.0000	0.0284	0.0450	0.0198	0.0502	0.0305
	Inference	0.0027	0.0015	0.0109	0.0042	0.0000	0.0009	0.0003	0.0002	0.0000	0.0294	0.0374	0.0192	0.0106	0.0108
Percentage Based (50%)	Symm	0.5000	0.5000	0.5000	0.1945	0.1875	0.3334	0.0000	0.0002	0.5000	0.5200	0.5000	0.4696	0.5000	0.5000
	Anti-Sym	0.0000	0.0005	0.0068	0.0041	0.0000	0.0002	0.0000	0.0002	0.0068	0.0305	0.0068	0.0045	0.0169	0.0106
	Inverse	0.0034	0.0038	0.0118	0.0053	0.0000	0.0003	0.0000	0.0003	0.0068	0.0286	0.0000	0.0005	0.0119	0.0080
Percentage Based (50%)	Inference	0.0041	0.0037	0.0048	0.0029	0.0007	0.0005	0.0004	0.0006	0.0088	0.0286	0.0136	0.0071	0.0147	0.0135

are performed on a local server with Intel Corporation Xeon E7 v4/Xeo CPU with 24 cores, 256 GB RAM, and GeForce GTX 1180 with 4 GPU cores.

## 5 Results

In this Section, we report the experiment results and discuss them. Tables 5 and 6 show the MRR and Hit@10 performance of the LP methods and the Tables 7, 8 and 9 report the AUC-PR results. Extended result sheets are available in [https://github.com/mlwin-de/relational\\_pattern\\_benchmarking](https://github.com/mlwin-de/relational_pattern_benchmarking).

**Inductive** As far as Inductive datasets are concerned, in both FB15K and WN18, GraIL and CompGCN outperform all other models due to their property of inductive bias-ness. Segregating far, all models performed well over the Inference dataset. CompGCN gave better performance on Inverse dataset with a 62.5 AUC-PR score whereas MDE did not perform well in all datasets from the inductive category. To sum up, the difference between Inductive and Transductive models could easily be noticed by this set of experiments.

An exception of performance for Symmetry dataset in Distmult for WN18 was observed with an increase of 17% from FB15K. Table 8 summarizes FB15K results and the green column of Table 7 summarizes the average AUC-PR on the Inductive setting.

All models show a low performance on the MRR and Hit metric for the Inductive datasets as in the Tables 5 and 6. However, the AUC-PR results show that some models can better distinguish positive and negative samples in an equal number of samples.

**Transductive** For evaluation with transductive setting of datasets, all models except CompGCN showed great improvement in link prediction task than their performance over other datasets. CompGCN was almost 10% - 15% less accurate as compared to other state of the art models. Symmetry dataset gives promising results with link prediction of more than 95% in almost all models. Furthermore, for WN18 dataset, GraIL and even CompGCN predicts 99% true triplets for the symmetry dataset. Overall TransE and RotatE are proved to be superior models for transductive datasets. Promising AUC-PR results for symmetry dataset by the model are displayed in Table 7.

Table 7: Mean AUC-PR performance of the LP methods.

Method	WN18				FB15K			
	Metric (AUC-PR)				Metric (AUC-PR)			
	Inductive	Transductive	Head/Tail Ratio	Percentage Based(50%)	Inductive	Transductive	Head/Tail Ratio	Percentage Based(50%)
DistMult	0.4745	0.6865	0.4454	0.5744	0.4473	0.6203	0.4449	0.5954
TransE	0.3296	0.6779	0.4028	0.5358	0.3877	0.7807	0.4491	0.5990
RotatE	0.3260	0.6554	0.3870	0.5208	0.3695	0.7341	0.3941	0.6160
MDE	0.3585	0.4689	0.4356	0.4560	0.3565	0.6437	0.4137	0.4998
GraIL	0.5000	0.6403	0.5000	0.5682	0.5000	0.7173	0.4973	0.6325
CompGCN	0.4502	0.6516	0.4920	0.5643	0.5000	0.7173	0.4973	0.6325
QuatE	0.9484	0.6194	0.5494	0.7817	0.9153	0.4526	0.5455	0.6929

Table 8: AUC-PR Results of Link Predictors on datasets extracted from FB15K.

Type of dataset	Dataset	FB15K					
		Metric (AUC-PR)					
		Distmult	TransE	RotatE	MDE	GraIL	CompGCN
Inductive	Symmetry	0.3646	0.3702	0.3712	0.3450	0.500	0.4497
	Anti-Symmetry	0.4412	0.3453	0.3138	0.3920	0.500	0.4968
	Inverse	0.4556	0.3251	0.3224	0.3230	0.500	0.6107
	Inference	0.5277	0.5103	0.4707	0.3660	0.500	0.5179
Transductive	Symmetry	0.9496	0.9854	0.9722	0.9690	0.9622	0.8999
	Anti-Symmetry	0.5199	0.6729	0.6075	0.4978	0.6416	0.5411
	Inverse	0.5872	0.7263	0.6568	0.4897	0.6670	0.5613
	Inference	0.4246	0.7382	0.6999	0.6182	0.5983	0.5979
Head/Tail Ratio	Symmetry	0.4976	0.4210	0.3875	0.4256	0.500	0.4458
	Anti-Symmetry	0.4436	0.4730	0.4113	0.4285	0.4973	0.4510
	Inverse	0.4452	0.4447	0.4086	0.4328	0.4827	0.4654
	Inference	0.3933	0.4579	0.3691	0.3679	0.5090	0.6160
Percentage based (50%)	Symmetry	0.8213	0.5840	0.7208	0.6678	0.8067	0.9042
	Anti-Symmetry	0.5215	0.5665	0.5467	0.4375	0.5713	0.5055
	Inverse	0.5571	0.6031	0.5530	0.4338	0.5917	0.4811
	Inference	0.4819	0.6422	0.6495	0.64601	0.6273	0.6685

**Semi-Inductive - Head Tail Ratio** Since this set of datasets is also inductive in its properties as all triplets are fully inductive with either of the entities unseen. The models behave exactly the same way as they work with the inductive datasets. All state-of-the-art models fail to perform under such settings except GraIL and CompGCN due to their inductive nature. Despite the fact that GraIL and CompGCN perform better than transductive bias model, they could not rank the triplets correctly with more than 0.52 AUC-PR. However, CompGCN shows some improvements for the Inference type datasets under the category with around 0.60 AUC-PR.

**Semi-Inductive - Percentage-Wise** The datasets are gathered on the principle of half datasets from the inductive category with 50% triplets of the transductive nature. Due to the reason, they perform well on the link prediction models. Distmult and CompGCN gives promising results of more than 85% on the symmetry datasets. TransE, due to its inability to infer symmetric patterns could not give better score but predicts inference patterns better than any other. As an average, we conclude that all models outperform over the transductive setting of datasets in link prediction task. The KGE models also give promising results over datasets with half inductive and half transductive type (Count-Based Dataset), evidently due to the fact of containing exactly 50% transductive triplets. Inductive datasets and head-tail inductive datasets still faces issues in link prediction task with better results on CompGCN and GraIL.

## 5.1 Discussion

**Inductive Evaluations** Since none of the three types of inductive experiments included unknown triples in the test phase, the models did not use any other information indicative of the identity of tested entities (e.g., their distance to other entities). Thus, the vector weights of the unknown entities are set randomly. Consequently, we conclude that the experimented models estimate the rank of the samples only by generating an embedding space such that the scores of negative triples are separable from the random values to some extent. In the following, we generally compare the performance of models over by datasets extracted from the two original knowledge graphs.



Table 9: AUC-PR Results of Link Predictors on datasets extracted from WN18.

Type of dataset	Dataset	WN18					
		Metric (AUC-PR)					
		DistMult	TransE	RotatE	MDE	GraIL	CompGCN
Inductive	Symmetry	0.5400	0.3600	0.3600	0.3654	0.500	0.3655
	Anti-Symmetry	0.4288	0.3256	0.3156	0.3696	0.500	0.4786
	Inverse	0.4832	0.3123	0.3089	0.3671	0.500	0.3646
	Inference	0.4459	0.3205	0.3197	0.3277	0.500	0.5920
Transductive	Symmetry	0.9934	0.9882	0.9939	0.9890	0.9940	0.9922
	Anti-Symmetry	0.5774	0.5340	0.5035	0.4341	0.5218	0.5395
	Inverse	0.5873	0.6119	0.5729	0.4798	0.5167	0.4624
	Inference	0.5879	0.5774	0.5512	0.4765	0.5286	0.6123
Head/Tail Ratio	Symmetry	0.4069	0.3450	0.3971	0.4434	0.500	0.4295
	Anti-Symmetry	0.4576	0.4306	0.3764	0.4437	0.500	0.4360
	Inverse	0.4629	0.4012	0.3835	0.4433	0.500	0.5145
	Inference	0.4549	0.4346	0.3913	0.4118	0.500	0.5879
Percentage based (50%)	Symmetry	0.8019	0.7225	0.7256	0.5224	0.7500	0.7580
	Anti-Symmetry	0.4845	0.4460	0.4641	0.4264	0.5087	0.4651
	Inverse	0.5323	0.5238	0.4407	0.4485	0.5068	0.4491
	Inference	0.4790	0.4508	0.4527	0.4267	0.5073	0.5505

Table 10: Evaluation results of GFA-NN model compared to MDE.

LP Model	Datasets									
	Metric (AUC-PR)									
	Transductive				Head/Tail Ratio		Percentage based (50%)			
	Symmetry	Anti-Symmetry	Inverse	Inference	Inverse	Inference	Symmetry	Anti-Symmetry	Inverse	Inference
MDE	0.9690	0.4978	0.4897	0.6182	0.4328	0.3679	0.6678	0.4375	0.4338	0.6460
GFA-NN (FB15K)	0.9750	0.7932	0.8150	0.7817	0.7817	0.7817	0.7191	0.6059	0.6976	0.7817
GFA-NN (WN18)	0.9937	0.7144	0.6245	0.6064	0.4687	0.4926	0.7224	0.4622	0.4673	0.4693

**FB15K** For FB15K, GraIL and CompGCN performs well on the inductive datasets with highest performance on inverse data set in CompGCN where as GraIL shows same performance on all datasets with inductive setting. CompGCN performs well on transductive datasets with almost 90% accuracy on Symmetry along with outstanding performance on all datasets with Count-based setting. GraIL shows exactly the same trend as CompGCN with similar accuracy. MDE performs well on transductive setting datasets with 96% accuracy on symmetry followed by count-based setting datasets. RotatE, TransE and DistMult showed the same performance pattern as MDE with highest performance in transductive setting and least in inductively set datasets. Table 8 describes the result of these evaluations.

**WN18** As far as WN18 is concerned, it shows similar pattern as the FB15K dataset. GraIL and CompGCN shows better results on inductive settings as compared to other models. Link prediction is better on WN18 datasets by 5-6% as compared to FB15K that evidently is due to the fact of less relations to be computed. Link prediction on Symmetry datasets from Transductive setting has around 0.99 AUC-PR for all models considered. Table 9 quotes results on WN18 Datasets.

## 5.2 Evaluation of GFA-NN model

Finally, we benchmark our datasets over a newer multi-objective optimization KGE similar to MDE, i.e., GFA-NN Sadeghi et al. (2021a), which considers the datasets’ graphical features to create embeddings. The node features and graph properties are calculated and stored in separate files for training purposes. The graphical feature information is not available to the model in the evaluation phase to make a fair comparison.

Comparison between the previous MDE and GFA-NN reported in Table 10 shows outstanding behavior of the newly introduced method. GFA-NN outperforms MDE by an increase of almost 8-10% in all datasets. The most significant improvement is observable in the head-tail inductive datasets, where the results show an improvement of more than 30% in the tested cases. For Instance, the Inference dataset of the setting improved from 36% in the AUC-PR score to 78%. Lastly, the results improved further in the Percentage-based setting, increasing about 10% in all cases. The inverse dataset of the category rose from 43% to 69% in the link prediction task.

Table 11: Ranking results of the LP models on the aggregate datasets.

Model	WN18RR			FB15k-237		
	MR	MRR	Hit@10	MR	MRR	Hit@10
QuatE	–	0.482	0.572	–	<b>0.366</b>	<b>0.556</b>
TransE	357	0.294	0.501	357	0.294	0.465
DistMult	5261	0.44	0.49	254	0.241	0.419
CompGCN	3533	0.479	0.546	197	0.355	0.535
RotatE	3340	0.476	0.571	177	0.338	0.533
MDE	3219	0.458	0.536	203	0.344	0.531
GFA-NN	3390	<b>0.486</b>	<b>0.575</b>	186	0.338	0.522

### 5.3 Comparison to Aggregate datasets

Table 11 shows the results of the LP methods on the aggregate datasets reported in Sadeghi et al. (2021a) and Sun et al. (2019).

This comparison of this table results to our pattern-specific experiments shows that while the difficulty of each individual relation pattern for each model influences the overall results, the ratio of each relation pattern also impacts their performance because the results are averaged in the MRR and Hit evaluations. For example, while CompGCN is not the best method in learning inverse relations, it is one of the most efficient methods in the FB15k-237 because FB15k-237 and WN18RR are missing the inverse relations from FB15K and WN18.

## 6 Conclusion

To sum up, we created several standard datasets with distinct relational patterns ranging from symmetry to inverse and different degrees of inductiveness. We evaluated several state-of-the-art Link Prediction models over them. We extended our research in working on link prediction models by this benchmarking approach and highlighting their working on different types of datasets. A unified evaluation strategy of AUC-PR measurement is incorporated into all link prediction models. We highlighted the datasets with improved performance on particular LP models. In order to support further research in the domain, we incorporated our benchmarking datasets to the BenchEmbedd for evaluation of a linked data life-cycle. The meaningful experiments outcome indicate that these datasets will foster the research on Link Prediction and KG embeddings.

## Acknowledgements

We acknowledge the support of the following projects: SPEAKER (BMW i FKZ 01MK20011A), JOSEPH (Fraunhofer Zukunftsstiftung), the EU projects Cleopatra (GA 812997), TAILOR (EU GA 952215), the BMBF project MLwin (01IS18050) and the BMBF excellence clusters ML2R (BmBF FKZ 01 15 18038 A/B/C) and ScaDS.AI (IS18026A-F).

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] We have both, datasets and Benchmarking
  - (b) Did you describe the limitations of your work? [N/A]
  - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [N/A] A summarized results is shown in a precise paper.

- (b) Did you include complete proofs of all theoretical results? [N/A] Our work includes no theoretical assumptions
3. If you ran experiments (e.g. for benchmarks)...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] A GitHub link is added which includes code, data and the instructions to reproduce both the dataset and the experiment results.
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Yes, we included them in the github repository in [https://github.com/mlwin-de/relational\\_pattern\\_benchmarking/blob/master/hyperparams.txt](https://github.com/mlwin-de/relational_pattern_benchmarking/blob/master/hyperparams.txt)
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A] The reported values are the average of 5 execution results.
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Yes. We included them in the Experiments setup Section.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [Yes] All reused studies and code and are cited
- (b) Did you mention the license of the assets? [Yes] We have added MIT License in the github repository in [https://github.com/mlwin-de/relational\\_pattern\\_benchmarking/blob/master/MIT-LICENSE.txt](https://github.com/mlwin-de/relational_pattern_benchmarking/blob/master/MIT-LICENSE.txt)
- (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] they are in github repo and also linked in the paper. We have added datasets.
- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] We use open sourced data only.
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] No, the data does not use any personal information
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We do not use crowd sourcing
- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## References

- Bollacker, K. D., Evans, C., Paritosh, P., Sturge, T., and Taylor, J. (2008). Freebase: a collaboratively created graph database for structuring human knowledge. In Wang, J. T., editor, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, pages 1247–1250. ACM.
- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Hao, Y., Zhang, Y., Liu, K., He, S., Liu, Z., Wu, H., and Zhao, J. (2017). An end-to-end model for question answering over knowledge base with cross-attention combining global knowledge. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 221–231, Vancouver, Canada. Association for Computational Linguistics.

- Hogan, A., Blomqvist, E., Cochez, M., d’Amato, C., de Melo, G., Gutierrez, C., Gayo, J. E. L., Kirrane, S., Neumaier, S., Polleres, A., Navigli, R., Ngomo, A.-C. N., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., and Zimmermann, A. (2021). Knowledge graphs.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. (2021). Open graph benchmark: Datasets for machine learning on graphs.
- Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P., Hellmann, S., Morsey, M., Van Kleef, P., Auer, S., and Bizer, C. (2014). Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal*, 6.
- Meilicke, C., Fink, M., Wang, Y., Ruffinelli, D., Gemulla, R., and Stuckenschmidt, H. (2018). Fine-grained evaluation of rule-and embedding-based systems for knowledge graph completion. In *International semantic web conference*, pages 3–20. Springer.
- Miller, G. A. (1995). Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41.
- Morris, C., Kriege, N. M., Bause, F., Kersting, K., Mutzel, P., and Neumann, M. (2020). Tudataset: A collection of benchmark datasets for learning with graphs.
- Rossi, A., Barbosa, D., Firmani, D., Matinata, A., and Merialdo, P. (2021). Knowledge graph embedding for link prediction. *ACM Transactions on Knowledge Discovery from Data*, 15(2):1–49.
- Sadeghi, A., Collarana, D., Graux, D., and Lehmann, J. (2021a). Embedding knowledge graphs attentive to positional and centrality qualities. In *European Conference on Machine Learning and Data Mining, ECML PKDD 2021*.
- Sadeghi, A., Graux, D., Yazdi, H. S., and Lehmann, J. (2020). MDE: Multiple distance embeddings for link prediction in knowledge graphs.
- Sadeghi, A., Shahini, X., Schmitz, M., and Lehmann, J. (2021b). Benchembedd: A fair benchmarking tool for knowledge graph embeddings. In *SEMANTiCS 2021 Demo Session*. <http://ceur-ws.org/>.
- Safavi, T. and Koutra, D. (2020). Codex: A comprehensive knowledge graph completion benchmark.
- Sarkar, A., Yang, Y., and Vihinen, M. (2020). Variation benchmark datasets: update, criteria, quality and applications. *Database*, 2020. baz117.
- Suchanek, F. M., Kasneci, G., and Weikum, G. (2007). Yago: A core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web, WWW ’07*, page 697–706, New York, NY, USA. Association for Computing Machinery.
- Sun, Z., Deng, Z.-H., Nie, J.-Y., and Tang, J. (2019). Rotate: Knowledge graph embedding by relational rotation in complex space.
- Teru, K. K., Denis, E., and Hamilton, W. L. (2020). Inductive relation prediction by subgraph reasoning.
- Vashishth, S., Sanyal, S., Nitin, V., and Talukdar, P. (2020). Composition-based multi-relational graph convolutional networks.
- Wang, H., Qin, K., Zakari, R. Y., Lu, G., and Yin, J. (2021). Deep neural network based relation extraction: An overview.
- Wang, Y., Ruffinelli, D., Gemulla, R., Broscheit, S., and Meilicke, C. (2019). On evaluating embedding models for knowledge base completion. In *Proceedings of the 4th Workshop on Representation Learning for NLP (Repl4NLP-2019)*, pages 104–112, Florence, Italy. Association for Computational Linguistics.
- Xiong, C., Power, R., and Callan, J. (2017). Explicit semantic ranking for academic search via knowledge graph embedding. In *Proceedings of the 26th International Conference on World Wide Web, WWW ’17*, page 1271–1279, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.

- Yang, B. and Mitchell, T. (2017). Leveraging knowledge bases in LSTMs for improving machine reading. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1436–1446, Vancouver, Canada. Association for Computational Linguistics.
- Yang, B., tau Yih, W., He, X., Gao, J., and Deng, L. (2015). Embedding entities and relations for learning and inference in knowledge bases.
- Yang, F., Yang, Z., and Cohen, W. W. (2017). Differentiable learning of logical rules for knowledge base reasoning. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2319–2328.
- Yousef, M., Jung, S., Showe, L., and Showe, M. (2008). Learning from positive examples when the negative class is undetermined- microrna gene identification. *Algorithms for molecular biology : AMB*, 3:2.
- Zeiler, M. D. (2012). ADADELTA: an adaptive learning rate method. *CoRR*, abs/1212.5701.
- Zhang, F., Yuan, N. J., Lian, D., Xie, X., and Ma, W.-Y. (2016). Collaborative knowledge base embedding for recommender systems. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 353–362, New York, NY, USA. Association for Computing Machinery.