

# Exploring the Limits of Few-Shot Link Prediction in Knowledge Graphs

Dora Jambor<sup>1,2</sup>, Komal Teru<sup>1,2</sup>, Joelle Pineau<sup>1,2,3</sup>, William L. Hamilton<sup>1,2</sup>

<sup>1</sup> School of Computer Science, McGill University, Canada

<sup>2</sup> Quebec AI Institute (Mila), Canada

<sup>3</sup> Facebook AI Research (FAIR)

{dora.jambor, komal.teru}@mail.mcgill.ca

{jpineau, wlh}@cs.mcgill.ca

## Abstract

Real-world knowledge graphs are often characterized by low-frequency relations—a challenge that has prompted an increasing interest in few-shot link prediction methods. These methods perform link prediction for a set of new relations, unseen during training, given only a few example facts of each relation at test time. In this work, we perform a systematic study on a spectrum of models derived by generalizing the current state of the art for few-shot link prediction, with the goal of probing the limits of learning in this few-shot setting. We find that a simple *zero-shot* baseline—which ignores any relation-specific information—achieves surprisingly strong performance. Moreover, experiments on carefully crafted synthetic datasets show that having only a few examples of a relation fundamentally limits models from using fine-grained structural information and only allows for exploiting the coarse-grained positional information of entities. Together, our findings challenge the implicit assumptions and inductive biases of prior work and highlight new directions for research in this area.

## 1 Introduction

A knowledge graph (KG) is a multi-relational graph that offers a structured way to organize facts about the world. Encoder-decoder approaches are commonly used to predict new facts from existing ones where entities and relations are embedded in a low-dimensional vector space via an encoder, to then score the likelihood of observing a new fact via a decoder (Nickel et al., 2015; Bordes et al., 2013; Trouillon et al., 2017; Dettmers et al., 2018).

It is well known that the performance of these methods can significantly drop when predicting for relations that are only observed in a few example facts. However, link prediction for these low-frequency relations is very important, as not only are these relations abundant in most knowledge graphs, they are also key for knowledge graph




Figure 1: The few-shot link prediction task

completion tasks where new relations may appear after model training.

To study this low-frequency regime, Xiong et al. (2018) created the Nell-One and Wiki-One benchmarks where the task is to predict new facts for a set of new relations at test time, where each relation is only observed a *few* times (as specified by some small fixed number  $K$ ). Previous approaches have shown promising results using metric-based (Xiong et al., 2018) and gradient based meta-learning techniques (Chen et al., 2019). However, we argue that these models are limited by the current task formulation to only exploit coarse-grained positional signals (i.e., nodes belonging to the same community) that are abundant in these benchmarks, rather than leveraging structural signals (e.g. transitivity, symmetry).

**Present work.** In this work, we take a critical take on current approaches for few-shot link prediction over knowledge graphs. We posit that current meta-learning based approaches benefit largely due to the positional signals in entities, rather than utilizing information about the low frequency relations. We corroborate these insights by conducting a systematic study on a spectrum of models with decreasing complexity. Interestingly, we find that a much simpler zero-shot variant of the state of the art—devoid of any meta-learning scheme—yields surprisingly competitive results, while not consuming any example facts about a relation. Motivated by these observations, we design a set of null models tailored to different learning signals a model might utilize to drive effective link predic-

tion. Empirically, we validate that these existing meta-learning models are ill-equipped to infer logical patterns about the few-shot relations. These findings bring forth the shortcomings of the current task formulation and raises new questions in both task and model design while highlighting new directions of research in few-shot link prediction.

## 2 Few-shot link prediction

### 2.1 Problem definition

The goal of few-shot link prediction is to predict missing links for a new relation by only observing  $K$  example triples of that relation (Figure 1). Following literature in few-shot classification (Ravi and Larochelle, 2016; Vinyals et al., 2016; Snell et al., 2017), we organize our dataset as a set of tasks, where a task corresponds to predicting links for a new relation. The set of tasks for training and testing are disjoint, with the added constraint that entities in the test tasks are a subset of the entities in the train tasks. Let  $\mathcal{V}$  denote the set of entities in the knowledge graph. For each new relation  $r_i$ , we then construct a support set  $\mathcal{S}_i = \{(h_k, r_i, t_k)\}_{k=1}^K$  containing  $K$  example entity pairs,  $h_k, t_k \in \mathcal{V}$ , connected by relation  $r_i$ , and a query set  $\mathcal{Q}_i = \{(h_j, r_i, ?)\}_{j=1}^J$  containing  $J$  query triples over entities in  $\mathcal{V}$ . As shown in Figure 1, the goal is then to learn how to extract knowledge from the support set such that we can predict the missing tail entities in the query set.

### 2.2 Overview of the framework

The foundation of our analyses focuses on a generalization of the current state-of-the-art gradient-based meta-learning approach (Chen et al., 2019). This approach follows the encoder-decoder paradigm of embedding-based knowledge graph completion methods (Hamilton et al., 2017), where the entities and relations are embedded in a low-dimensional vector space and the embeddings are used to predict the likelihood of a given triple.

**Encoder functions.** The key idea in few-shot learning is to transfer knowledge from support set to query set by learning a function  $\text{RelLearner} : \mathcal{S}_i \mapsto \mathbb{R}^d$ . This maps a support set  $\mathcal{S}_i$ , which characterizes the relation  $r_i$ , to a low dimensional embedding via an encoder function  $\mathbb{E} : \mathcal{V} \mapsto \mathbb{R}^d$

$$\mathbf{r}_i = \text{RelLearner}(\{\{\mathbb{E}(h_k), \mathbb{E}(t_k)\}_{k=1}^K\}). \quad (1)$$

The  $\text{RelLearner}$  function can vary from a simple MLP (Hastie et al., 2009) to more complicated

recurrent architectures (Rumelhart et al., 1985; Jordan, 1997; Hochreiter and Schmidhuber, 1997). Further, the entity encoder  $\mathbb{E}$  can vary from TransE-style embeddings (Bordes et al., 2013; Sun et al., 2019) to a graph neural network (Schlichtkrull et al., 2018) that explicitly leverages the neighborhood information around entities.

**Decoder and loss function.** A decoder function ingests the embeddings of the entities  $\mathbf{h}$ ,  $\mathbf{t}$  and of the relation  $\mathbf{r}$  to score the likelihood of a given triple  $(h, r, t)$ . Using a simple TransE decoder (Bordes et al., 2013), it is then optimized to score positive triples higher than negative triples using a contrastive loss  $\mathcal{L}$  (Dyer, 2014). In the few-shot setting we compute the support set loss  $\mathcal{L}(\mathcal{S}_i)$ , and the final query set loss  $\mathcal{L}(\mathcal{Q}_i)$ , which are used to update the model parameters (Chen et al., 2019).

**Meta-gradient update.** Instead of directly using the relation embedding  $\mathbf{r}_i$  from Equation (1) to compute the final query loss  $\mathcal{L}(\mathcal{Q}_i)$ , we first make an update on the relation embedding using the gradient of the support set loss  $\mathcal{L}(\mathcal{S}_i)$

$$\mathbf{r}'_i = \mathbf{r}_i - \eta \nabla_{\mathbf{r}_i} \mathcal{L}(\mathcal{S}_i) \quad (2)$$

where  $\eta$  denotes the learning rate. This update encourages  $\mathbf{r}_i$  to be such that it effectively predicts the support set triples via minimizing  $\mathcal{L}(\mathcal{S}_i)$ .

### 2.3 Baselines

Our objective is to probe how much models leverage the support set to perform the query task. To this end, we perform a systematic study of different model variants, where each falls into the general framework described in Section 2.2.

**MetaR** follows Chen et al. (2019), where the  $\text{RelLearner}$  is defined as a 2-layer MLP (Hastie et al., 2009) over the support set entity embeddings. The encoder  $\mathbb{E}$  simply maps each entity to a fixed learnable vector as in Bordes et al. (2013).

**SharedEmbed** skips Equation (1), and instead sets  $\mathbf{r}_i = \mathbf{r}_g$ , where  $\mathbf{r}_g$  is a single learnable embedding shared across all relations. We propose this modification to measure the effect of representing all relations by the same embedding  $\mathbf{r}_g$ , where the only information from the support set comes via the gradient update in Equation (2).

**ZeroShot** further removes the meta-gradient update in Equation (2) and lets  $\mathbf{r}'_i = \mathbf{r}_g$ . This effectively reduces the model to perform zero-shot link prediction on the relation’s query set without any relation-specific information from the support set.

		MRR		Hits@10		Hits@5		Hits@1	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
<b>Nell-One</b>	MetaR	0.294	0.323	0.464	0.500	0.398	0.426	0.201	0.230
	SharedEmbed	0.276	0.311	0.454	0.495	0.382	0.420	0.173	0.205
	ZeroShot	0.199	0.219	0.342	0.365	0.283	0.303	0.116	0.136
	R-GCN	0.216	0.267	0.412	0.464	0.316	0.366	0.120	0.172
<b>Wiki-One*</b>	MetaR	0.325	0.326	0.448	0.408	0.408	0.367	0.258	0.280
	SharedEmbed	0.290	0.311	0.399	0.415	0.348	0.378	0.238	0.254
	ZeroShot	0.279	0.289	0.361	0.367	0.337	0.341	0.234	0.246
	R-GCN	0.126	0.137	0.178	0.237	0.130	0.152	0.101	0.104

Table 1: Average metrics on Nell-One and Wiki-One few-shot link prediction tasks. \* For Wiki-One we used pre-trained embeddings, similar to Chen et al. (2019).

**R-GCN** uses the same `RelLearner` as in MetaR, with the exception that support set entity embeddings are learned via a multi-relational graph neural network, R-GCN (Schlichtkrull et al., 2018), instead of a TransE-style embeddings. R-GCN learns entity representation via aggregating the 2-hop neighbors of a given entity. With this model we probe the extent to which injecting structural bias into entity representations can influence performance in link prediction.

### 3 Null Models

In order to probe and understand the performance of different models, we introduce two *null models*, which are used to generate synthetic data that satisfy certain properties. Motivated by recent literature on position versus structure-aware methods in relational learning (You et al., 2019; Srinivasan and Ribeiro, 2020), we test the models’ ability to learn from two key sources of information: *structural information* and *positional information*. In the context of knowledge graphs, *structural information* corresponds to the fine-grained relational semantics. These are the logical patterns that are extracted by state-of-the-art rule induction systems, such as RuleN (Meilicke et al., 2018).

On the other hand, *positional information* corresponds to the coarse-grained community structure of the nodes in the graph. In other words, two nodes are said to be positionally ‘close’ in the graph, i.e., if they belong to the same community (Newman, 2018).

#### 3.1 Structural Null Models

The first type of null models contains synthetic relations that satisfy simple logical properties. For the sake of exposition, we focus on two simple logical patterns: symmetry and transitivity. For the purposes of all synthetic data generation, we

only consider the largest connected component of respective datasets, denoted  $\mathcal{G}^L$ .

**Synthetic symmetric relations.** To generate  $2N$  edges connected by a symmetric relation  $r_s^*$ , we repeat the following steps  $N$  times.:

1. Uniformly sample a pair of unique entities— $h_i, t_i$ —from all the entities in  $\mathcal{G}^L$ .
2. Add two edges— $((h_i, r_s^*, t_i)), ((t_i, r_s^*, h_i))$  to the set of synthetic symmetric edges.

**Synthetic transitive relations.** To sample  $3N$  edges connected by a transitive relation  $r_t^*$ , we generate 3 edges at a time. In particular, we repeat the following steps  $N$  times:

1. Uniformly sample 3 unique entities— $e_1, e_2$ , and  $e_3$ —from all the entities in  $\mathcal{G}^L$ .
2. Add three edges— $(e_1, r_t^*, e_2), (e_2, r_t^*, e_3), (e_1, r_t^*, e_3)$ —to our collection.

#### 3.2 Positional Null Models

The second type of null models focuses on generating synthetic relations that depend on the underlying community structure in the graph. We call these relations *positional* because they depend on the relative global position of the entities, rather than on local structural properties.

We first cluster the largest connected component  $\mathcal{G}^L$  into  $K$  communities using a standard algorithm originally proposed by Blondel et al. (2008). Let  $\{C_i\}_{i=1}^K$  denote the set of communities generated, where each community is a set of entities from  $\mathcal{G}^L$ . To generate  $N$  synthetic edges for a positional relation  $r_p$ , we repeat the following steps  $N$  times:

1. Uniformly sample a community index  $i$  from the set  $\{1, \dots, K\}$ .
2. Uniformly sample two unique entities  $h, t$  from community  $C_i$ ; add  $(h, r_p^*, t)$  to the set.




Figure 2: Average Hits@10 on synthetically generated relations using our proposed null models.




Figure 3: Pearson’s R between MRR and the log frequency of support set entities in training graph. We observe strong correlation for Nell-One, but not for Wiki-One. For more details on this, see Appendix C.

## 4 Experiments

We followed the same experimental setup as in Chen et al. (2019), as described in Appendix A. We conducted our experiments on the Nell-One and Wiki-One datasets<sup>1</sup>. For more details on these benchmarks, we refer the reader to Table 1 in Xiong et al. (2018). Similar to earlier work, we report MRR, Hits@1, Hits@5 and Hits@10 on our test relations, using a type-constrained candidate set.

### 4.1 Results and Analysis

**Experiments on Real Data.** As shown in Table 1, for Nell-One, we find that SharedEmbed model yields competitive performance to MetaR, with Hits@10 of 45.4% and 49.5%, as compared to MetaR’s Hits@10 of 46.4% and 50.0% for 1 and 5-shot, respectively. The same observation holds for Wiki-One, where SharedEmbed yields 39.9% and 41.5% Hits@10, compared to MetaR’s Hits@10 of 44.8% and 40.8%, for 1 and 5-shot, respectively.

It is surprising how competitive SharedEmbed is, given that the only relation-specific information the model gets to observe comes via the meta-gradient update in Equation (2). In fact, we find that even in absence of this gradient signal, i.e., without *any* relation-specific information, ZeroShot performs relatively good, with Hits@10 of 34.2% and 36.5% on Nell-One, and 36.1% and 36.7% on Wiki-One.

The nontrivial performance of these simple models suggests that such models may exploit some easily accessible positional signals around entities,

without the need to learn meaningful representations for relations. In fact, Figure 3 shows a high correlation between performance and the degrees of entities in the support set for Nell-One. We reconcile this observation by noting that as models observe more signals about entities, they start relying less on the support set, and thus on the relation representations. Furthermore, contrary to our expectation, even when we equip models with structural biases, as done via an R-GCN, they do not yield better results.

**Null Model Experiments.** We probed the above trained models on the synthetically generated test tasks following the procedure discussed in Section 3. As shown in Figure 2, we find a consistent trend for these models to yield higher performance on tasks that rely on positional signals, as compared to tasks that require logical inference.

Indeed, in the current task formulation, where we are given a support set of  $K$  randomly sampled examples, it is unlikely that logically consistent patterns will be captured in the  $K$ -shot examples. For example, seeing conclusive evidence of transitivity when only given a small random sample of tuples is highly unlikely. In fact, as we show in Appendix B, one provably cannot learn certain logical patterns for some values of  $K$  in the  $K$ -shot setting.

## 5 Conclusion

We conducted a systematic study of various models to probe their limits in performing few-shot link prediction. Our experiments on both synthetic and real data show that the current task formulation encourages models to mainly rely on positional information around entities, rather than leveraging logical signals about relations. In fact, we empirically show that having only  $K$  examples of a relation fundamentally limits the types of logical patterns that can be learned. We argue that a future direction in few-shot link prediction should allow for a more careful construction of the support set, to scaffold the use of logical patterns in few-shot learning.

<sup>1</sup>Datasets can be downloaded under [this link](#).

## Acknowledgments

The authors would like to thank Priyesh Vijayan, Joey Bose, Lu Liu and other members of the RLLab and Mila for their invaluable feedback and useful discussions on the early drafts. Furthermore, the authors are grateful for the anonymous reviewers for their comments on the first draft of the paper. This research was supported in part by Canada CIFAR AI Chairs held by Prof. Pineau and Prof. Hamilton, as well as gift grants from Microsoft Research and Samsung AI.

## References

- Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, pages 2787–2795.
- Mingyang Chen, Wen Zhang, Wei Zhang, Qiang Chen, and Huajun Chen. 2019. Meta relational learning for few-shot link prediction in knowledge graphs. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4217–4226. Association for Computational Linguistics.
- Tim Dettmers, Minervini Pasquale, Stenertorp Pontus, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Chris Dyer. 2014. Notes on noise contrastive estimation and negative sampling. *arXiv preprint arXiv:1410.8251*.
- William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.*
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Michael I Jordan. 1997. Serial order: A parallel distributed processing approach. In *Advances in psychology*, volume 121, pages 471–495. Elsevier.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Christian Meilicke, Manuel Fink, Yanjie Wang, Daniel Ruffinelli, Rainer Gemulla, and Heiner Stuckenschmidt. 2018. Fine-grained evaluation of rule- and embedding-based systems for knowledge graph completion. In *International Semantic Web Conference*, pages 3–20. Springer.
- Mark Newman. 2018. *Networks*. Oxford university press.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2015. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33.
- Sachin Ravi and Hugo Larochelle. 2016. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1985. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087.
- Balasubramaniam Srinivasan and Bruno Ribeiro. 2020. On the equivalence between positional node embeddings and structural graph representations. In *International Conference on Learning Representations*.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*.
- Théo Trouillon, Christopher R Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. 2017. Knowledge graph completion via complex tensor factorization. *The Journal of Machine Learning Research*, 18(1):4735–4772.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638.
- Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2018. One-shot relational learning for knowledge graphs. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1980–1990. Association for Computational Linguistics.

Jiaxuan You, Rex Ying, and Jure Leskovec. 2019.  
Position-aware graph neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 72*.

## A Hyperparameters

As discussed, we followed the experimental setup described in [Chen et al. \(2019\)](#). We used the Adam optimizer ([Kingma and Ba, 2015](#)) with a learning rate of 0.001, using 1 to 3 ratio of positive to negative samples. During training, we used 3 queries per task on each dataset.

We adapted the batch size to be 1024, and the number of queries to test on to be 3, based on their open-sourced codebase. These hyperparameters yielded the best performing models. Similarly, we also used the same train/validation/test relation splits of 51:5:11 and 133:16:34 for Nell-One and Wiki-One respectively.

For our R-GCN model, we considered a range of [5, 10, 20] as the number of neighbors to sample for each message passing step, and [2, 4] as the number of basis. Furthermore, we used 2 layers in the R-GCN. Finally, we used 50 and 20 as the R-GCN hidden layer dimension for Nell-One and Wiki-One, respectively. These hyperparameters were partly followed from [Schlichtkrull et al. \(2018\)](#), and were decided upon consideration for our available compute infrastructure.

Our models were trained on a single Nvidia 1080Ti GPU, and each model training took between 13-18 hours depending on the model and dataset settings.

## B Logical inference in the few-shot domain

Figure 4 shows five example support sets to demonstrate that inferring logical properties such as symmetry and transitivity requires a minimum number of carefully designed K-shot examples.

## C Entity Frequency Analysis

Figure 5 shows the 100 highest-degree entities out of all 4,838,244 entities in the Wiki-One knowledge graph. We find that the median degree of entities is 1, and the highest degree is 227,390, which connects to 4.69% of the total graph. We suspect that these high-degree entities, so-called hub nodes, may add noise to the embeddings of support set entities. This could in turn affect performance and explain why we do not observe a strong correlation between the degrees of support set entities and performance in Wiki-One.

K-shot	Example support sets of varying K-shots.	Query triple	Answer
3	[(Bob, coworker_of, Alice), (Alice, coworker_of, Ella), (Ella, coworker_of, Bob)]	(Stella, coworker_of, ?)	N/A
4	[(Bob, coworker_of, Alice), (Alice, coworker_of, Ella), (Ella, coworker_of, Bob), (Stella, coworker_of, Alice)]	(Stella, coworker_of, ?)	Ella, Bob
5	[(Bob, coworker_of, Alice), (Alice, coworker_of, Ella), (Ella, coworker_of, Bob), (Larry, coworker_of, Jim), (Jim, coworker_of, Laura)]	(Laura, coworker_of, ?)	Larry
2	[(Liz, friend_of, Rob), (Rob, friend_of, Liz)]	(Peter, friend_of, ?)	N/A
3	[(Liz, friend_of, Rob), (Rob, friend_of, Liz), (Barbara, friend_of, Peter)]	(Peter, friend_of, ?)	Barbara

Figure 4: Limits of logical inference in the few-shot domain.




Figure 5: Entity degree distribution of Wiki-One