

Learning Embeddings from Knowledge Graphs With Numeric Edge Attributes

Sumit Pai, Luca Costabello

Accenture Labs

{sumit.pai, luca.costabello}@accenture.com

Abstract

Numeric values associated to edges of a knowledge graph have been used to represent uncertainty, edge importance, and even out-of-band knowledge in a growing number of scenarios, ranging from genetic data to social networks. Nevertheless, traditional knowledge graph embedding models are not designed to capture such information, to the detriment of predictive power. We propose a novel method that injects numeric edge attributes into the scoring layer of a traditional knowledge graph embedding architecture. Experiments with publicly available numeric-enriched knowledge graphs show that our method outperforms traditional numeric-unaware baselines as well as the recent UKGE model.

1 Introduction

Knowledge graphs are graph-based knowledge bases whose facts are modeled as labeled, directed edges between entities. Whether it is a social network, a bioinformatics dataset, or retail purchase data, modelling knowledge as a graph lets organizations capture patterns that would otherwise be overlooked. Research led to broad-scope graphs such as DBpedia [Auer *et al.*, 2007], WordNet, and YAGO [Suchanek *et al.*, 2007]. Countless domain-specific knowledge graphs have also been published on the web, giving birth to the so-called Web of Data [Bizer *et al.*, 2011].

Knowledge graph embeddings (KGE) are a family of graph representation learning methods that learn vector representations of nodes and edges of a knowledge graph. They are applied to graph completion, knowledge discovery, entity resolution, and link-based clustering, just to cite a few [Nickel *et al.*, 2016].

In *multimodal* knowledge graphs, entity nodes are associated to attributes from different data modalities, such as numbers, text, images. In this paper we deal with a specific flavor of multimodal knowledge graphs, that is graphs with numeric-enriched edges, either at predicate-level (i.e. a number assigned to each predicate *type*), or specific to each triple (this latter case is shown in Figure 1). Numeric-enriched triples have a prominent role in a growing number of applicative scenarios, from protein networks to workforce knowledge bases. Such numeric values may repre-

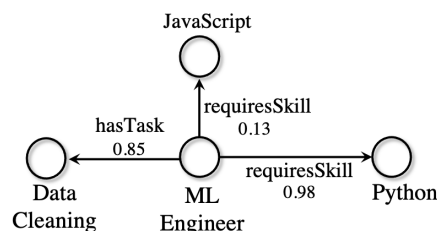


Figure 1: A Knowledge graph with numeric attributes associated to triples.

sent edge uncertainty (i.e. uncertain knowledge graphs) as in ConceptNet [Speer *et al.*, 2017]; trust in the automated or semi-automated data extraction workflow used to build the graph [Mitchell *et al.*, 2018]; out-of-band knowledge from wet lab experiments [Szklarczyk *et al.*, 2016]; edge importance or link strength.

A number of works in knowledge graph representation learning literature support multimodal information and leverage numeric values associated to node entities to generate better embeddings for enhanced link prediction [Garcia-Duran and Niepert, 2017; Kristiadi *et al.*, 2019; Wu and Wang, 2018]. Nevertheless, such models are not designed to learn from numeric values associated to *edges* of a knowledge graph. With the notable exception of [Chen *et al.*, 2019], that however is designed for uncertain graphs only, supporting numeric is to date still an under-researched direction.

In this work, we focus on the task of predicting probability estimates of missing links in knowledge graphs with numeric-enhanced triples. We claim that a model must take into account such numeric literals. Regardless of their semantics, we operate under the assumption that such numeric values intensify or mitigate the probability of existence of a link. That implies treating triples with low-numeric values as “pseudo-negatives”. Traditional embedding models put at the same level low-valued and high-valued triples. However, in many real-world scenarios it is important to make such distinction (protein interactions [Szklarczyk *et al.*, 2016]).

We propose FocusE, an add-on layer for knowledge graph embeddings to enhance link prediction with edge-related numeric literals. FocusE works with any existing KGE model that adopts the standard negatives generation protocol [Bor-

des *et al.*, 2013]. We use edge numeric literals to modulate the margin between the scores of true triples and their corresponding negative corruptions. Inspired by Focal Loss [Lin *et al.*, 2017] that aims at sparser hard examples by modulating the loss function, we leverage numeric literals to “focus” traditional KGE models on triples with higher numeric values.

Experiments show that models trained using FocusE outperform numeric-unaware baselines, in particular in discriminating triples with high-numeric attributes from those associated to low values.

2 Related Work

Knowledge Graph Embeddings. Although a comprehensive survey is out of the scope of this work (recent surveys provide a good coverage of the landscape [Bianchi *et al.*, 2020]), it is worth listing the most popular knowledge graph embedding models proposed to date. TransE [Bordes *et al.*, 2013] is the forerunner of distance-based models, and inspired a number of models commonly referred to as TransX. The symmetric bilinear-diagonal model DistMult [Yang *et al.*, 2015] paved the way for its asymmetric evolutions in the complex space, ComplEx [Trouillon *et al.*, 2016] and RotatE [Sun *et al.*, 2019]. Some models such as RESCAL [Nickel *et al.*, 2011], TuckER [Balažević *et al.*, 2019], and Simple [Kazemi and Poole, 2018] rely on different tensor decomposition techniques. Models such as ConvE [Dettmers *et al.*, 2018] or ConvKB [Nguyen *et al.*, 2018] leverage convolutional layers. Attention is used by [Nathani *et al.*, 2019].

Numeric-aware models (node attributes). None of the models listed above leverage numeric attributes of any kind. However, a number of recent works support multimodal knowledge graphs and learn from numeric values associated to node entities. LiteralE enriches node embeddings with numeric information before scoring the triples [Kristiadi *et al.*, 2019]. KBRLN combines latent, relational and numeric features using product of experts model [Garcia-Duran and Niepert, 2017]. TransEA learns a vanilla structural model using TransE scoring, and an attribute model for attributed triples, using regression over the attribute values, which is jointly trained [Wu and Wang, 2018]. Such models are not designed to learn from numeric values associated to *edges*.

Numeric-aware models (edge attributes). To the best of our knowledge¹, the only work designed to work with numeric-aware edges is UKGE [Chen *et al.*, 2019]. UKGE generates confidence scores for known triples by squashing numeric values in the $[0 - 1]$ interval. It then uses probabilistic soft logic [Kimmig *et al.*, 2012] to predict probability estimates for unseen triples, by jointly training a model to regress over the confidence values. A limitation of this approach is

¹We limit to Knowledge Graph Embeddings literature and knowledge graphs, i.e. directed, labeled, multi graphs. Other numeric-aware models such as graph neural networks that operate on different data modalities are out of the scope of this paper (they are either designed for tasks other than link prediction, most of them do not support multi-relational graphs, and cannot currently be applied at the scale KGE models operate).

that out-of-band logical rules are required as additional input. It is also worth noting that UKGE design rationale aims at supporting uncertain knowledge graphs, i.e. graphs whose edge numeric values represent uncertainty. In this paper we aim at supporting generic numeric values, regardless of their semantics (uncertainty, strength, importance, etc). Moreover, we claim that triples with high numeric values should have higher impact, and this should be explicitly modeled in the scoring or loss computation.

3 Preliminaries

Knowledge Graph. A knowledge graph $\mathcal{G} = \{(s, p, o)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ is a set of triples $t = (s, p, o)$ each including a subject $s \in \mathcal{E}$, a predicate $p \in \mathcal{R}$, and an object $o \in \mathcal{E}$. \mathcal{E} and \mathcal{R} are the sets of all entities and relation types of \mathcal{G} .

Knowledge Graph with numeric-enriched triples. In a numeric-enriched knowledge graph \mathcal{G} , each triple is assigned a numeric attribute $w \in \mathbb{R}$, leading to $\mathcal{G} = \{t = (s, p, o, w)\}$. It is worth noting that we are not tied to a specific numeric attribute semantics, as these numbers may encode importance, uncertainty, strength, etc. For example, Figure 1 assumes that numeric values indicate the importance of a link. The triple (MLEngineer, requiresSkill, Python, 0.98) is therefore more important than (MLEngineer, requiresSkill, JavaScript, 0.13).

Note w can be either defined at a predicate level or at triple level. In this paper, we assume that w is triple-specific.

Knowledge Graph Embedding Models. Knowledge graph embedding models (KGE) are neural architectures designed to predict missing links between entities. KGE encode both entities \mathcal{E} and relations \mathcal{R} into low-dimensional, continuous vectors $\in \mathbb{R}^k$ (i.e. the embeddings). Knowledge graph embeddings are learned by training a neural architecture over a training knowledge graph: an input layer feeds training triples to an embedding lookup layer that retrieves embeddings of entities and relations.

A scoring layer $f(t)$ assigns plausibility scores to each triple. The scoring layer is designed to assign high scores to positive triples and low scores to negative corruptions. Most of literature differs in the design rationale of $f(t)$. For example, the scoring function of TransE [Bordes *et al.*, 2013] computes a similarity between the embedding of the subject e_s translated by the embedding of the predicate e_p and the embedding of the object e_o .

Corruptions are synthetic negative triples generated by a corruption generation layer that follows the protocol proposed in [Bordes *et al.*, 2013]: we define a corruption of t as $t^- = (s, p, o')$ or $t^- = (s', p, o)$ where s', o' are respectively subject or object corruptions (i.e. other entities randomly selected from \mathcal{E}). We generate synthetic negatives by corrupting one side of the triple at a time to comply with the local closed world assumption [Nickel *et al.*, 2016].

Finally, a loss layer optimizes the embeddings by maximizing the margin between positive triples t and corruptions t^- . In other words, the goal of the optimization procedure is learning optimal embeddings, such that at inference time the scoring function $f(t)$ assigns high scores to triples likely to be correct and low scores to triples unlikely to be true.

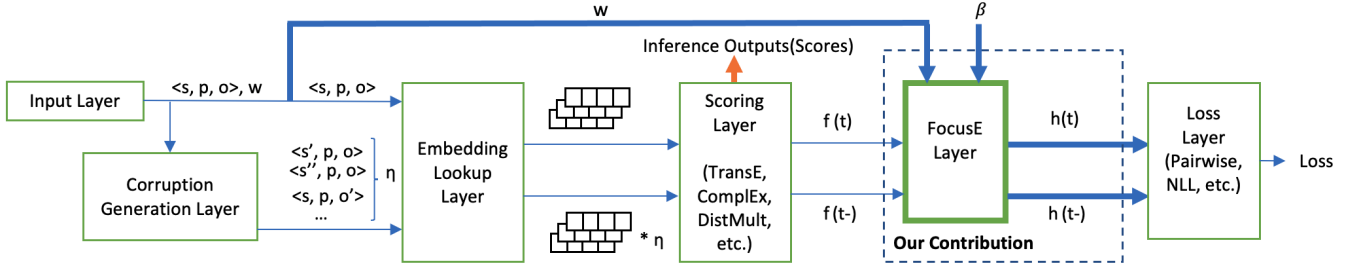


Figure 2: A Knowledge Graph Embedding model architecture enhanced with FocusE. The add-on acts as an intermediate layer between the traditional scoring layer and the loss.

Link Prediction. The task of predicting unseen triples in knowledge graphs is formalized in literature as a learning to rank problem, where the objective is learning a scoring function $f(t = (s, p, o)) : \mathcal{E} \times \mathcal{R} \times \mathcal{E} \rightarrow \mathbb{R}$ that given an input triple $t = (s, p, o)$ assigns a score $f(t) \in \mathbb{R}$ proportional to the likelihood that the fact t is true. Such predictions are ranked against predictions from synthetic corruptions, to gauge how well the model tells positives from negatives.

Link Prediction with numeric-enriched triples. In this paper we predict probability estimates for unseen numeric-enhanced triples $t = (s, p, o, w)$. The task is formalized as the same learning to rank problem described above for conventional link prediction.

4 FocusE

We present FocusE, an add-on layer for knowledge graph embedding architectures designed for link prediction with numeric-enriched triples. FocusE takes into account numeric literals associated to each link. Regardless of their semantics, we operate under the assumption that numeric values intensify or mitigate the probability of existence of a link. For example, given numeric values w in the $[0 - 1]$ range, we assume that high values identify triples with higher chances of being true, low scores single out weak or unlikely relations, and $w = 0$ triples are considered negative samples.

FocusE consists in a plug-in layer that fits between the scoring and loss layers of a conventional KGE method and it is designed to be used during training (Figure 2). Unlike traditional architectures, before feeding the scoring layer to a loss function, we modulate its output based on numeric values associated to triples, to obtain “focused” scores. We leverage numeric values associated to triples so that during training the model focuses on triples with higher numeric values. We want our model to learn from training triples with high numeric values, and at the same time use edge numeric values to maximise the margin between scores assigned to true triples and those assigned to their corruptions. This increases the loss of the model and helps it focus on triples with higher values. Our contribution is described below in detail.

Let $t = (s, p, o)$ be a positive triple and $w > 0$ its numeric-value. We define a corruption of t as $t^- = (s, p, o')$ or $t^- = (s', p, o)$, where s', o' are respectively subject or object corruptions.

Let $f(t)$ be the scoring function of a KGE model. In the case of TransE [Bordes *et al.*, 2013] this is:

$$f(t) = -\|\mathbf{e}_s + \mathbf{r}_p - \mathbf{e}_o\|_n \quad (1)$$

where \mathbf{e}_s , \mathbf{r}_p and \mathbf{e}_o are the embeddings of the subject s , predicate p , and object o .

We use a softplus non-linearity σ to make sure the scores returned by $f(t)$ are greater or equal to zero, without introducing excessive distortion:

$$g(t) = \sigma(f(t)) = \ln(1 + e^{f(t)}) \geq 0 \quad (2)$$

To take into account the effect of numeric values associated to triples, we define a modulating factor $\alpha \in \mathbb{R}$ which is responsible to strike a balance between the influence of the structure of the graph and the impact of the numeric values associated to each triple:

$$\alpha = \begin{cases} \beta + (1 - w)(1 - \beta) & \text{if } t \\ \beta + w(1 - \beta) & \text{if } t^- \end{cases} \quad (3)$$

where $\beta \in [0, 1]$ is the *structural influence*, an hyperparameter that modulates the influence of graph topology, and $w \in \mathbb{R}$ is the numeric value associated to the *positive* triple t . β is used to re-weigh the triple value w . If $\beta = 0$ the original numeric values w are used. If $\beta = 1$, numeric values w are ignored and the model is equivalent to a conventional KGE architecture. Note that positive and negative triples are assigned different α equations. This is done to lower the margin between the scores for triples and their respective corruptions when the triple numeric value is high.

Finally, the FocusE layer $h(t)$ is defined as:

$$h(t) = \alpha g(t) \quad (4)$$

Putting all together, the FocusE layer $h(t)$ is then used in the loss function² L . This is a modified, more numerically stable version of the negative log-likelihood of normalized softmax scores proposed in [Kadlec, 2017]:

$$L = - \sum_{t^+, t^-} \log \frac{e^{h(t^+)}}{e^{h(t^+)} + e^{h(t^-)}} \quad (5)$$

²Note that the FocusE layer described in Eq.4 is compatible with other losses used in KGE literature (e.g. pairwise, negative log-likelihood, etc.)

As seen in equation 5, the modulation between structural influence and numeric values increases the margin between high and low-valued triples. Hence the model learns to focus on triples with higher numeric attributes.

During training, we decay the structural influence β from 1 to 0. This is controlled by the hyperparameter λ (decay). Initially the model gives equal importance to all training triples. When β decays to zero (linearly, over λ epochs), the model exclusively relies on numeric attributes (w). Section 5.3 shows empirical evidence of the effectiveness of λ .

5 Experiments

We assess the predictive power of FocusE on the link prediction task with numeric-enriched triples. Experiments show that FocusE outperforms conventional KGE models and its closest direct competitor UKGE [Chen *et al.*, 2019] in discriminating low-valued triples from high-valued ones.

Datasets

We experiment with three publicly available benchmark datasets originally proposed by [Chen *et al.*, 2019]. Triples are associated to numeric values w interpreted as uncertainty values. We also introduce a fourth dataset, where numeric values w encode the *importance* of each link. All datasets include triple-specific w values. Table 1 shows the statistics of all the datasets used for the experiments.

- **CN15K** [Chen *et al.*, 2019]. A subset of ConceptNet [Speer *et al.*, 2016], a common sense knowledge graph built to represent general human knowledge. Numeric values on triples represent uncertainty.
- **NL27K** [Chen *et al.*, 2019]. A subset of the Never-Ending Language Learning (NELL) [Mitchell *et al.*, 2018] dataset, which collects data from web pages. Numeric values on triples represent link uncertainty.
- **PP15K** [Chen *et al.*, 2019]. Knowledge graph of protein-protein interactions [Szklarczyk *et al.*, 2016]. Numeric values represent the confidence of the link based on existing scientific literature evidence.
- **O*NET20K**³. We introduce a subset of O*NET⁴, a dataset that includes job descriptions, skills and labeled, binary relations between such concepts. Each triple is labeled with a numeric value that indicates the *importance* of that link. Unlike the other datasets, we built a test set that includes a single predicate type that connects jobs to skills, under the assumption that we are interested in predicting that type of links only. This is done to mock up single-target link prediction tasks adopted in applied knowledge discovery scenarios.

Implementation Details and Baselines. FocusE and all baselines are implemented with the AmpliGraph library [Costabello *et al.*, 2019] version 1.4.0, using TensorFlow 1.15.2 and Python 3.7. Code and experiments are available at <https://github.com/Accenture/AmpliGraph>. We experiment with three popular KGE baselines: TransE, DistMult,

	O*NET20K	CN15K	NL27K	PP15K
Training	461,932	204,984	149,100	230,929
Validation [‡]	138	3532	8161	1940
Test (top 10%)*	200	1,929	1,402	2,172
Test (bottom 10%)*	200	1,929	1,402	2,172
Entities	20,643	15,000	27,221	5,000
Relations	19	36	404	7

Table 1: Datasets used in experiments. (*) test sets include either top-10% valued triples or bottom 10%. All w are normalized, so that $w \in [0 - 1]$. (‡) validation sets only include high-valued triples where $w \geq 0.8$.

ComplEx. For each baseline and for FocusE, we carried out extensive grid search, over the following ranges of hyperparameter values: embedding dimensionality $k = [200 - 600]$, with a step of 100; baseline losses={negative log-likelihood, multiclass-NLL, self-adversarial}; synthetic negatives ratio $\eta = \{5, 10, 20, 30\}$; learning rate= $\{1e-3, 5e-3, 1e-4\}$; epochs= $[100 - 800]$, step of 100; L3 regularizer, with weight $\gamma = \{1e-1, 1e-2, 1e-3\}$. For FocusE we also tuned the decay $\lambda = [100 - 800]$, with increments of 100.

For experiments with UKGE we modified⁵ the original codebase⁶ provided by the authors to generate learning to rank metrics required by our experimental protocol. We used the best hyperparameter configuration proposed by the authors. Note we do not evaluate UKGE on O*NET20K, since the model requires additional logical rules which are not available for this dataset. All experiments were run under Ubuntu 16.04 on an Intel Xeon Gold 6142, 64 GB, equipped with a Tesla V100 16GB.

Evaluation protocol. We adopt the standard evaluation protocol described by [Bordes *et al.*, 2013] to all datasets described above. We predict whether each triple $t = (s, p, o) \in \mathcal{T}$ is a positive fact, where \mathcal{T} is a disjoint held-out test set that includes only positives triples. We cast the problem as a learning-to-rank task: for each $t = (s, p, o) \in \mathcal{T}$, we generate synthetic negatives $t^- \in \mathcal{N}_t$ by corrupting one side of the triple at a time (i.e. either the subject or the object). We predict a score for each t and all its negatives $t^- \in \mathcal{N}_t$. We then rank the only positive t against all the negatives \mathcal{N}_t ⁷. We report learning to rank metrics such as mean rank (MR), mean reciprocal rank (MRR), and Hits at n (where $n = 1, 10$) by filtering out spurious ground truth positives from the list of generated corruptions (i.e. “filtered” metrics).

⁵The original UKGE codebase does not perform a fair rank computation: the authors assign rank=1 when t and t^- have the same score, whereas we assign a lower rank, as per agreed upon practice in the community. For example, if we have $\eta = 1,000$ corruptions t^- which are all assigned the same score as t , we assign t a rank=1,000 whereas UKGE ranks it first. To guarantee a fair comparison, we aligned UKGE to our procedure.

⁶<https://github.com/stasl0217/UKGE>

⁷As for standard protocol, we used distinct entities twice, leading to CN15K= \sim 30k, NL27K= \sim 54k, PP15K=10k, O*NET20K= \sim 40k synthetic negatives (Table 1).

³<https://docs.ampligraph.org/en/latest/ampligraph.datasets.html>

⁴<https://www.onetonline.org/>

MRR (bottom 10%)				
	O*NET20K	CN15K	NL27K	PP15K
TransE	.13	.05	.36	.26
DistMult	.15	.05	.83	.97
ComplEx	.12	.08	.81	.95
UKGE _{rect}	-	.02	.28	.00*
UKGE _{logi}	-	.02	.36	.56
FocusE TransE	.14	.04	.26	.08
DistMult	.15	.05	.70	.72
ComplEx	.13	.08	.53	.49

Table 2: Predictive power of FocusE on the 10% low-valued triples. Lower = better. Best results in bold. (*) UKGE_{rect} fails to produce actionable results on PPI5K.

5.1 Predicting High-Valued Links

First, we assess how well FocusE predicts triples associated to *high* numeric values. This is an important task in many applicative scenarios: high numeric values usually imply low uncertainty, high strength or importance, thus leading to more valuable newly-discovered knowledge. Table 3 reports, for each dataset, the best mean rank, MRR, and Hits@{1,10} computed over test sets that include only triples associated to the top-10% numeric attributes. Our goal is learning to assign low ranks to test triples (i.e. rank=1 being the best outcome).

Results show that FocusE brings better or very similar MRR to traditional, numeric-unaware baselines: on O*NET20K, FocusE increases MRR for all models, and it outperforms the best baseline by 14 base points (MRR marginally differs on CN15K, NL27K, and PPI5K).

Experiments show that FocusE outperforms UKGE: on CN15K, MRR is 15 base points higher, 19 points higher on NL27K, and up to 30 points better for PPI15K (UKGE_{rect} fails to provide actionable results⁸ on PPI15K). Besides, UKGE requires additional external rules, hence the absence of results for O*NET20K. FocusE achieves better predictive power, without requiring additional out-of-band rules.

5.2 Discriminating High and Low-Valued Links

The primary goal of FocusE is achieving a clear separation between scores assigned to high and low-valued triples. To assess how well FocusE can tell high-valued links from low-valued, we look at two complementary tasks: how the model performs on test triples with high values (top 10%) and on triples with low values (bottom 10%). We collect the respective MRR results and we compute the difference Δ MRR. The higher the difference, the better. Note this task differs from what traditionally done by KGE models, which are designed to tell positives from synthetic negatives only, thus not being able to discriminate between low and high-valued triples.

Such prior art shortcoming is evident from results in Table 4. Models trained using FocusE have a much wider Δ MRR compared to traditional baselines. Using link numeric values of training triples, FocusE better discriminates triples with higher values in the test set from the ones with lower values (see also Table 2 for details on the bottom 10% links). On

⁸Due to a problem in the UKGE codebase, UKGE predictive power is MRR=0, hence "not actionable".

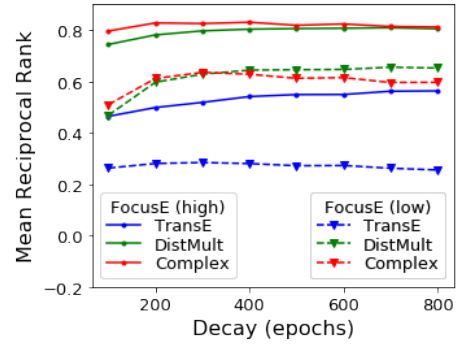


Figure 3: Impact of λ , the decay of structural influence β on model performance over n-epochs for NL27k dataset.

O*NET20K, NL27K and PPI5K, FocusE models perform exceptionally better, with Δ MRR being consistently higher than traditional KGE baselines. FocusE also outperforms UKGE by a large margin (75%) on CN15K. It provides a large margin of 40 base points on PPI5K as well, i.e. a 411% increase (UKGE_{rect} does not provide actionable results).

5.3 Decay Impact

The structural influence β affects the impact of numeric values w . If $\beta = 1$, FocusE falls back to the underlying KGE model. If $\beta = 0$, the original numeric values w have maximum impact throughout the entire training. During training we decay β linearly, until it reaches 0.

In this experiment we vary the decay λ , while freezing the other hyperparameters to ($k = 300$, $\eta = 30$, NLL of normalized softmax scores loss (Equation 5), Adam optimizer with learning rate $1e-3$, L3 regularizer, with weight $\gamma = 1e-2$). We choose different decay epoch values $\lambda = \{0, 100, 200, 300, 400, 500, 600, 700, 800\}$. We report changes in MRR on NL27K's top 10% test set by experimenting on TransE, DistMult, and ComplEx. Figure 3 shows that performance improves if λ increases. In most cases model performance saturates when $\lambda > 400$ epochs.

6 Conclusion

We show that by plugging an additional layer we can make a conventional KGE architecture aware of numeric values associated to triples. This leads to models that better discriminate high-valued and low-valued triples, regardless of the semantics of the numeric attributes, and without requiring additional out-of-band rules (unlike UKGE).

Future work will investigate the capability of predicting numeric values associated to unseen triples. We will also extend our approach to support multiple numeric attributes associated to the same triple.

References

- [Auer *et al.*, 2007] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In *The semantic web*, pages 722–735. Springer, 2007.

		O*NET20K (top 10%)				CN15K (top 10%)			
		MR	MRR	Hits@		MR	MRR	Hits@	
				1	10			1	10
	TransE	<u>8</u>	.37	.00	.88	<u>623</u>	.19	.03	.47
	DistMult	13	<u>.49</u>	<u>.33</u>	.82	942	.22	.15	.37
	ComplEx	26	.32	.19	.59	863	<u>.28</u>	.22	.40
	UKGE _{rect}	-	-	-	-	984	.14	.05	.32
	UKGE _{logi}	-	-	-	-	2214	.13	.08	.22
FocusE	TransE	4	.41	.0	.94	438	.18	.02	.47
	DistMult	102	.63	.52	<u>.85</u>	1552	.24	.17	.37
	ComplEx	9	.46	.23	.83	840	.29	<u>.21</u>	<u>.45</u>
		NL27K (top 10%)				PP15k (top 10%)			
		MR	MRR	Hits@		MR	MRR	Hits@	
				1	10			1	10
	TransE	<u>91</u>	.56	.38	.86	6	.42	.00	.94
	DistMult	130	.83	.77	<u>.93</u>	6	.97	.96	.99
	ComplEx	90	.87	.82	.94	10	.97	<u>.95</u>	.99
	UKGE _{rect}	169	.61	.50	.80	4943	.00	.00	.00
	UKGE _{logi}	299	.65	.56	.80	2	.65	.37	.99
FocusE	TransE	90	.57	.37	.88	4	.43	< .01	<u>.96</u>
	DistMult	140	.82	.77	.92	16	<u>.96</u>	.93	.99
	ComplEx	224	<u>.84</u>	<u>.79</u>	.92	20	.95	.92	.99

Table 3: Predicting high-valued links: ranking metrics computed on test sets that consist in the top-valued 10% triples. Filtered metrics. Best results in bold, second best underlined.

	O*NET20K			CN15K		
	MRR (top 10%)	MRR (bottom 10%)	Δ MRR	MRR (top 10%)	MRR (bottom 10%)	Δ MRR
TransE	.37	.13	.24	.19	.05	.14
FocusE TransE	.41	.14	.27 (+13%)	.18	.04	.14
DistMult	.49	.15	.34	.22	.05	.17
FocusE DistMult	.63	.15	.48 (+41%)	.24	.05	.19 (+12%)
ComplEx	.32	.12	.20	.28	.08	.20
FocusE ComplEx	.46	.13	.33 (+65%)	.29	.08	.21 (+5%)
UKGE _{rect}	-	-	-	.14	.02	.12
UKGE _{logi}	-	-	-	.13	.02	.11
FocusE ComplEx	.46	.13	.33	.29	.08	.21 (+75%)
	NL27K			PP15K		
	MRR (top 10%)	MRR (bottom 10%)	Δ MRR	MRR (top 10%)	MRR (bottom 10%)	Δ MRR
TransE	.56	.36	.20	.42	.26	.16
FocusE TransE	.57	.26	.31 (+55%)	.43	.08	.35 (+119%)
DistMult	.83	.83	< .01	.97	.97	< .01
FocusE DistMult	.82	.70	.12 (+1,100%)	.96	.72	.24 (+2,300%)
ComplEx	.87	.81	.06	.97	.95	.02
FocusE ComplEx	.84	.53	.31 (+417%)	.95	.49	.46 (+2,200%)
UKGE _{rect}	.61	.28	.33	.00	.00	.00
UKGE _{logi}	.65	.36	.29	.65	.56	.09
FocusE ComplEx	.84	.53	.31 (-6%)	.95	.49	.46 (+411%)

Table 4: High-valued vs low-valued triples discriminative power: FocusE brings larger differences in MRR across the board, showing better capabilities at correctly ranking the top-10% high-valued triples vs the bottom 10%. Best results in bold.

- [Balažević *et al.*, 2019] Ivana Balažević, Carl Allen, and Timothy M Hospedales. Tucker: Tensor factorization for knowledge graph completion. *arXiv preprint arXiv:1901.09590*, 2019.
- [Bianchi *et al.*, 2020] Federico Bianchi, Gaetano Rossiello, Luca Costabello, Matteo Palmonari, and Pasquale Minervini. Knowledge Graph Embeddings and Explainable AI. *arXiv preprint arXiv:2004.14843*, 2020.
- [Bizer *et al.*, 2011] Christian Bizer, Tom Heath, and Tim Berners-Lee. Linked data: The story so far. In *Semantic services, interoperability and web applications: emerging concepts*, pages 205–227. IGI Global, 2011.
- [Bordes *et al.*, 2013] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *NIPS*, pages 2787–2795, 2013.
- [Chen *et al.*, 2019] Xuelu Chen, Muhao Chen, Weijia Shi, Yizhou Sun, and Carlo Zaniolo. Embedding uncertain knowledge graphs. In *AAAI*, volume 33, pages 3363–3370, 2019.
- [Costabello *et al.*, 2019] Luca Costabello, Sumit Pai, Chan Le Van, Rory McGrath, Nicholas McCarthy, and Pedro Tabacof. AmpliGraph: a Library for Representation Learning on Knowledge Graphs, 2019.
- [Dettmers *et al.*, 2018] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *AAAI*, 2018.
- [Garcia-Duran and Niepert, 2017] Alberto Garcia-Duran and Mathias Niepert. KBLRN: End-to-end learning of knowledge base representations with latent, relational, and numerical features. *arXiv preprint arXiv:1709.04676*, 2017.
- [Kadlec, 2017] Rudolf et al. Kadlec. Knowledge base completion: Baselines strike back. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 69–74, 2017.
- [Kazemi and Poole, 2018] Seyed Mehran Kazemi and David Poole. Simple embedding for link prediction in knowledge graphs. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *NeurIPS*, pages 4284–4295. Curran Associates, Inc., 2018.
- [Kimmig *et al.*, 2012] Angelika Kimmig, Stephen Bach, Matthias Broecheler, Bert Huang, and Lise Getoor. A short introduction to probabilistic soft logic. In *Proceedings of the NIPS Workshop on Probabilistic Programming: Foundations and Applications*, pages 1–4, 2012.
- [Kristiadi *et al.*, 2019] Agustinus Kristiadi, Mohammad Asif Khan, Denis Lukovnikov, Jens Lehmann, and Asja Fischer. Incorporating literals into knowledge graph embeddings. In *International Semantic Web Conference*, pages 347–363. Springer, 2019.
- [Lin *et al.*, 2017] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaïming He, and Piotr Dollár. Focal loss for dense object detection. *CoRR*, abs/1708.02002, 2017.
- [Mitchell *et al.*, 2018] Tom Mitchell, William Cohen, Hruschka, et al. Never-ending learning. *Communications of the ACM*, 61(5):103–115, 2018.
- [Nathani *et al.*, 2019] Deepak Nathani, Jatin Chauhan, Charu Sharma, and Manohar Kaul. Learning attention-based embeddings for relation prediction in knowledge graphs. *arXiv preprint arXiv:1906.01195*, 2019.
- [Nguyen *et al.*, 2018] Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. A novel embedding model for knowledge base completion based on convolutional neural network. In *NAACL*, pages 327–333, 2018.
- [Nickel *et al.*, 2011] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A three-way model for collective learning on multi-relational data. In *ICML*, 2011.
- [Nickel *et al.*, 2016] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *Procs of the IEEE*, 104(1):11–33, 2016.
- [Speer *et al.*, 2016] Robyn Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. *CoRR*, abs/1612.03975, 2016.
- [Speer *et al.*, 2017] Robert Speer, Joshua Chin, and Catherine Havasi. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. In *AAAI*, pages 4444–4451, 2017.
- [Suchanek *et al.*, 2007] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: a core of semantic knowledge. In *Procs of WWW*, pages 697–706. ACM, 2007.
- [Sun *et al.*, 2019] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *ICLR*, 2019.
- [Szkłarczyk *et al.*, 2016] Damian Szklarczyk, John H Morris, et al. The string database in 2017: quality-controlled protein–protein association networks, made broadly accessible. *Nucleic acids research*, page gkw937, 2016.
- [Trouillon *et al.*, 2016] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *ICML*, pages 2071–2080, 2016.
- [Wu and Wang, 2018] Yanrong Wu and Zhichun Wang. Knowledge graph embedding with numeric attributes of entities. In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 132–136, 2018.
- [Yang *et al.*, 2015] Bishan Yang, Scott Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In *ICLR*, 2015.