
Stay Positive: Knowledge Graph Embedding Without Negative Sampling

Ainaz Hajimoradlou¹ Seyed Mehran Kazemi²

Abstract

Knowledge graphs (KGs) are typically incomplete and we often wish to infer new facts given the existing ones. This can be thought of as a binary classification problem; we aim to predict if new facts are true or false. Unfortunately, we generally only have positive examples (the known facts) but we also need negative ones to train a classifier. To resolve this, it is usual to generate negative examples using a *negative sampling* strategy. However, this can produce false negatives which may reduce performance, is computationally expensive, and does not produce calibrated classification probabilities. In this paper, we propose a training procedure that obviates the need for negative sampling by adding a novel regularization term to the loss function. Our results for two relational embedding models (DistMult and Simple) show the merit of our proposal both in terms of performance and speed.

1. Introduction

Relational embedding models have proved effective for link prediction in knowledge graphs (KGs) (Nickel et al., 2016; Nguyen, 2017; Wang et al., 2017; Kazemi et al., 2020). These models learn a mapping from each entity $e \in \mathcal{E}$ and each relation $r \in \mathcal{R}$ to a representation in a latent space such that a pre-defined score function ϕ operating in that latent space can make a prediction about whether a triple (h, r, t) is true or not based on the representations for h , r , and t .

Ideally, if a KG contained a set of facts and a set of falsehoods (corresponding to positive and negative examples), one could formulate link prediction as a binary classification problem. However, a particular problem in training link prediction models is that KGs only contain facts (i.e., positive examples). To address this problem, link prediction models resort to negative sampling techniques to generate a

Table 1. The percentage of time spent on negative sampling and on the rest of the operations for each training epoch of Simple and ComplEx on Fb15k. The results are averaged over 100 epochs. Note that besides the negative sampling column, the other column is also affected by negative sampling.

Model	Library	Negative Sampling	Rest
Simple	PyTorch	69.3%	30.7%
ComplEx	Theano	62.3%	37.7%

set of triples that are expected to have a low chance of being true and use those triples as negative examples.

Current negative sampling strategies have been criticized for several reasons including generating large numbers of false-negatives, generating negative examples that are far from the model’s decision boundary and produce near-zero loss, and making the models produce uncalibrated probabilities (Wang et al., 2018; Shan et al., 2018). Moreover, as evidenced in Table 1 for two sample models (ComplEx (Trouillon et al., 2016) and Simple (Kazemi & Poole, 2018)¹), a large portion of training time is typically devoted to generating negative samples. Using negative examples during training either requires a substantial increase in the size of mini-batches and consequently in the required amount of GPU memory or requires increasing the number of mini-batches which slows down training. This can cause issues for real-world KGs that are orders of magnitude larger than the benchmark datasets (Balkir et al., 2019). Nevertheless, negative sampling remains the mainstream approach primarily due to the lack of better alternatives.

In this paper, we obviate the need for negative sampling through proposing a one-class classification framework where we train models using only the positive examples and a novel regularization loss. As an orthogonal contribution, we identify issues with WN18RR (Dettmers et al., 2018) and FB15k-237 (Toutanova & Chen, 2015), two well-established benchmarks for link prediction in KGs, and create cleaned versions that resolve the identified issue.

¹BorealisAI, Vancouver, Canada ²BorealisAI, Montreal, Canada. Correspondence to: Ainaz Hajimoradlou <ainaz.hajimoradlou@borealisai.com>.

¹The timings are for the FB15k dataset (Bollacker et al., 2008) using the best hyperparameters reported for these models. For ComplEx, we used the implementation provided by the authors at <https://github.com/ttrouill/complEx>. For Simple, we used the fast version available at <https://github.com/baharefatemi/Simple>.

2. Background and Notation

For a set \mathcal{E} of entities and \mathcal{R} of relations, let $\zeta_{\mathcal{E},\mathcal{R}}$ represent the set of all triples of the form (h, r, t) such that $h, t \in \mathcal{E}$ and $r \in \mathcal{R}$ where the triple is a fact. A *knowledge graph* (KG) $\mathcal{G}_{\mathcal{E},\mathcal{R}} \subset \zeta_{\mathcal{E},\mathcal{R}}$ contains a subset of the facts in $\zeta_{\mathcal{E},\mathcal{R}}$. *KG completion* is the problem of inferring $\zeta_{\mathcal{E},\mathcal{R}}$ from $\mathcal{G}_{\mathcal{E},\mathcal{R}}$. A relational embedding model $\mathcal{M}(\theta)$ maps each entity and each relation to a hidden representation known as *embedding* and defines a function $\phi_{\mathcal{M}(\theta)} : \mathcal{E} \times \mathcal{R} \times \mathcal{E} \rightarrow \mathbb{R}$ from the head and tail entities and the relation in a triple to a plausibility score. To convert the score into a probability, one can take the Sigmoid of the score.

Bilinear Models: are a successful class of embedding models. A bilinear model learns a vector embedding $\mathbf{z}_e \in \mathbb{R}^d$ for each entity $e \in \mathcal{E}$ and a matrix embedding $\mathbf{Z}_r \in \mathbb{R}^{d \times d}$ for each relation $r \in \mathcal{R}$, where d is the embedding dimension. The score function for a triple (h, r, t) is then defined as $\phi_{\mathcal{M}(\theta)}(h, r, t) = \mathbf{z}_h^T \mathbf{Z}_r \mathbf{z}_t$. Bilinear models mainly differ in the restrictions they impose on the \mathbf{Z}_r matrices.

Training: Let $\mathcal{G}_{\mathcal{E},\mathcal{R}}$ be a KG and $\mathcal{M}(\theta)$ be a model. The parameters θ are typically learned by minimizing a loss function similar to the following:

$$\mathcal{L}(\theta) = \sum_{(h,r,t) \in \mathcal{G}_{\mathcal{E},\mathcal{R}}} \left(\mathcal{L}^+(\phi_{\mathcal{M}(\theta)}(h, r, t)) + \sum_{(h',r',t') \in \text{Neg}(h, r, t; n)} \mathcal{L}^-(\phi_{\mathcal{M}(\theta)}(h', r', t')) \right) \quad (1)$$

where $\text{Neg}(h, r, t; n)$ is a negative sampling function that produces n negative examples for the given positive example (h, r, t) and \mathcal{L}^+ and \mathcal{L}^- are the loss functions for positive and negative triples respectively. The hyperparameter n is called the *negative ratio*. A common choice for \mathcal{L}^+ and \mathcal{L}^- , also used in our experiments, is the negative log-likelihood: $\mathcal{L}^l(\phi_{\mathcal{M}(\theta)}(h, r, t)) = \text{softplus}(-l * \phi_{\mathcal{M}(\theta)}(h, r, t))$.

3. An Alternative to Negative Sampling

We define our loss function as:

$$\mathcal{L}(\theta) = \sum_{(h,r,t) \in \mathcal{G}_{\mathcal{E},\mathcal{R}}} \mathcal{L}^+(\phi_{\mathcal{M}(\theta)}(h, r, t)) + \lambda \mathcal{L}^{sp}(\phi_{\mathcal{M}(\theta)}) \quad (2)$$

where the first term aims at penalizing low scores for the positive examples in $\mathcal{G}_{\mathcal{E},\mathcal{R}}$, $\mathcal{L}^{sp}(\phi_{\mathcal{M}(\theta)})$ aims at preventing the model from generating high scores for all triples as well as giving the model a prior probability about the correctness of triples (*sp* stands for *stay positive*), and λ is a hyperparameter. We define the regularization term as:

$$\mathcal{L}^{sp}(\phi_{\mathcal{M}(\theta)}) = \left\| \sum_{h \in \mathcal{E}} \sum_{r \in \mathcal{R}} \sum_{t \in \mathcal{E}} \phi_{\mathcal{M}(\theta)}(h, r, t) - \psi |\mathcal{E}|^2 |\mathcal{R}| \right\|_p \quad (3)$$

\mathcal{L}^{sp} encourages the sum of the model scores for all triples to be close to $\psi |\mathcal{E}|^2 |\mathcal{R}|$ (or equivalently the average to be close to ψ)². By encouraging the average score to be around ψ , for a randomly constructed triple (h, r, t) we have $\mathbb{E}[\phi_{\mathcal{M}(\theta)}(h, r, t)] \approx \psi$. Therefore, $\sigma(\psi)$ can be considered as the prior probability of the model about the correctness of triples (recall that ϕ provides a score and $\sigma(\phi)$ provides a probability). If besides the positive samples in the KG we also know (an approximation of) the ratio of positive to negative triples (corresponding to the prior probability of the correctness of triples), we can set ψ accordingly. Otherwise, we consider ψ as a hyperparameter.

In practice, we observed that optimizing the sum of all scores in Equation (3) to be close to zero is easier than optimizing it to be close to a non-zero number. To take this observation into account, instead of optimizing for a model $\mathcal{M}(\theta)$, we optimize another model $\mathcal{M}'(\theta)$ where $\phi_{\mathcal{M}'(\theta)} = \phi_{\mathcal{M}(\theta)} + \psi$. Then, Equation (2) can be re-written for $\mathcal{M}'(\theta)$ as $\mathcal{L}(\theta) = \sum_{(h,r,t) \in \mathcal{G}_{\mathcal{E},\mathcal{R}}} \mathcal{L}^+(\phi_{\mathcal{M}(\theta)}(h, r, t) + \psi) + \lambda \mathcal{L}^{sp}(\phi_{\mathcal{M}'(\theta)})$ and the regularization term can be re-written as $\mathcal{L}^{sp}(\phi_{\mathcal{M}(\theta)}) = \left\| \sum_{h \in \mathcal{E}} \sum_{r \in \mathcal{R}} \sum_{t \in \mathcal{E}} \phi_{\mathcal{M}(\theta)}(h, r, t) \right\|_p$.

Regularization vs closed world assumption (CWA):

CWA (corresponding to assuming any triple that is not in the KG is false) is a specific case of our model with ψ set to $\frac{|\mathcal{G}_{\mathcal{E},\mathcal{R}}|}{|\mathcal{E}|^2 |\mathcal{R}|}$, where $|\mathcal{G}_{\mathcal{E},\mathcal{R}}|$ corresponds to the number of facts in the KG. But our regularization is more general in that as ψ becomes larger, the CWA becomes more relaxed. Unlike negative sampling, training with our regularization term never asks the model to provide a negative score for a specific triple. Instead, it provides the model with the flexibility of assigning any score to any triple as long as the mean of the scores matches with ψ .

Time Complexity: While a naive calculation of Equation (3) requires $O(|\mathcal{E}|^2 |\mathcal{R}|)$ score computations we prove that for bilinear models it can be computed efficiently.

Theorem 1. *Let \mathcal{M} be a bilinear model with embedding dimension d where for each relation, only $d \leq \alpha \leq d^2$ of the elements in the embedding matrix can be non-zero. For a knowledge graph $\mathcal{G}_{\mathcal{E},\mathcal{R}}$, the sum of scores of \mathcal{M} for all possible triples can be computed in $O(|\mathcal{E}|d + |\mathcal{R}|\alpha)$ as $(\sum_{e \in \mathcal{E}} \mathbf{z}_e^T)(\sum_{r \in \mathcal{R}} \mathbf{Z}_r)(\sum_{e \in \mathcal{E}} \mathbf{z}_e)$.*

Corollary 1. *Since $\alpha \in O(d^2)$ for RESCAL, for a knowledge graph $\mathcal{G}_{\mathcal{E},\mathcal{R}}$, the sum of the RESCAL scores for all possible triples can be computed in $O(|\mathcal{E}|d + |\mathcal{R}|d^2)$.*

Corollary 1 can be extended to the models that use a combination of bilinear score functions such as neural tensor networks (Socher et al., 2013).

²In each batch, we compute the regularization term only for entities and relations in the batch; computing it for all entities and relations did not offer an improvement in our experiments.

Corollary 2. Since $\alpha \in O(d)$ in *DistMult*, *ANALOGY*, *Simple*, *ComplEx*, *HolE* (and several other bilinear models), for a knowledge graph $\mathcal{G}_{\mathcal{E},\mathcal{R}}$, the sum of the model scores for all possible triples for any of the above models can be computed in $O((|\mathcal{E}| + |\mathcal{R}|)d)$.

For a KG $\mathcal{G}_{\mathcal{E},\mathcal{R}}$, let β represent the average number of triples each entity in \mathcal{E} appears in. Let n represent the negative ratio. Let \mathcal{M} be a bilinear model with embedding dimension d where for each relation, only α of the elements in the embedding matrix can be non-zero. Assume $\alpha \geq d$ (this is true for all models mentioned so far) and $|\mathcal{E}| > |\mathcal{R}|$ (in general and in each batch). The following theorems establish the time complexity of a single training epoch using Equation 1 and Equation 2.

Theorem 2. The time complexity of a single training epoch using Equation 1 is $O(\beta(1+n)|\mathcal{E}|\alpha)$.

Theorem 3. The time complexity of a single training epoch using Equation 2 is $O(\beta|\mathcal{E}|\alpha)$.

Theorems 2 and 3 indicate that the time complexity of a single epoch when training with our regularization term is $(1+n)$ times less than training with negative sampling.

A practical consideration: In Equation (3), a few scores may become extremely large and affect the total sum in an undesired way. For this reason, in our experiments, for the models trained with the regularization in Equation (3), we restrict the elements of the embedding vectors to be between -1 and 1 (by applying an element-wise Tanh function) resulting in $\phi_{\mathcal{M}} \in (-d, d)$ where d is the embedding dimension. Then we multiply the model score by $\frac{I}{d}$, where I is a hyperparameter, leading to $\phi_{\mathcal{M}} \in (-I, I)$.

4. Experiments and Results

We experiment with *DistMult* and *Simple* as two representative bilinear models and provide head-to-head comparisons between their original training and the proposed training in this paper. We use *spDistMult* and *spSimple* to refer to *DistMult* and *Simple* trained using our regularization term, where *sp* stands for *Stay Positive*.

Datasets: Two established benchmarks for link prediction in KGs are *WN18RR* (Dettmers et al., 2018) and *FB15k-237* (Toutanova & Chen, 2015). We noticed that a large number of entities that appear in validation and test sets of these two datasets do not appear in any training triples. For example, *WN18RR* contains 209 unseen entities in the test set. This is problematic because the predictions for these entities is mainly a function of initialization. As an independent contribution, we clean these two datasets further by removing the triples from validation and test sets involving unseen entities. We call the new datasets *WN18AM* and *FB15k-237AM*. *WN18AM* contains 40,559 entities, 11 relations,

86,835 train triples, 2,824 validation triples and 2,924 test triples. *FB15k-237AM* contains 14,505 entities, 237 relations, 272,115 train triples, 17,526 validation triples and 20,438 test triples. We use the two new datasets as well as two other standard benchmarks that contain negative samples in their test sets so we can compare model probability calibrations. The two datasets are *WN11* and *FB13* both introduced in (Socher et al., 2013).

Classification and calibration: From the obtained results on *WN11* and *FB13* reported in Table 2 (the columns under “Original”), one can see clear boosts for *spDistMult* and *spSimple* in terms of NLL (up to 77% improvement), Brier score (up to 49%), and also AUC (up to 20%). In Figure 1, we bin the output probabilities of different models on the test sets of *WN11* and *FB13* and estimate a probability density function using kernel density estimation (Terrell & Scott, 1992). It can be viewed that the output probabilities of *DistMult* and *Simple* are mostly close to zero and squeezed in one bin; this explains their poor performance in terms of NLL and Brier score. These plots also match the plots of Tabacof & Costabello (2020) for *ComplEx*. For *spDistMult* and *spSimple*, however, it can be viewed that the output probabilities are more spread across bins and centered around 0.5. This explains (in part) the superior performance of *spDistMult* and *spSimple* in terms of NLL and Brier score compared to *DistMult* and *Simple*.

Post-calibration: Tabacof & Costabello (2020) show empirically that post-calibration can substantially improve the output probabilities of the relational embedding models. Here, we examine how post-calibration affects our models and baselines. We use Platt scaling (Platt et al., 1999) for post-calibration. In Table 2 (“Post-Calibrated” columns), we observe that post-calibration improves both the NLL and Brier score of the four models. Since *DistMult* and *Simple* are originally highly uncalibrated, they are substantially impacted by post-calibration. For *spDistMult* and *spSimple*, however, the impact is much smaller because they already produce better-calibrated outputs. Moreover, it can be viewed that even after post-calibration, in most cases *DistMult* and *Simple* perform inferior compared to *spDistMult* and *spSimple* without post-calibration. Note that such post-calibrations are only possible if one has access to ground truth positive and negative examples in a validation set.

Link prediction: Table 2 represents the results obtained on *WN18AM* and *FB15k-237AM*. As can be viewed, without using any negative examples, *spDistMult* and *spSimple* achieve comparable performance compared to *DistMult* and *Simple*, showing that the regularization term can act as a replacement for negative samples. For further analysis, we measured the performance of the models and baselines when applying no dropout. We observed that the filtered MRR

Table 2. Results on four datasets. The columns under “Original” demonstrate the results of the model without any post-processing and the columns under “Post-Calibrated” demonstrate the results after post-calibrating the models using Platt scaling.

Model	WN11					FB13					WN18AM		FB15k-237AM	
	Original			Post-Calibrated		Original			Post-Calibrated		MRR	Hit@1	MRR	Hit@1
	NLL	Brier	AUC	NLL	Brier	NLL	Brier	AUC	NLL	Brier				
DistMult	2.670	0.430	0.668	0.675	0.241	1.572	0.393	0.695	0.630	0.219	0.440	0.410	0.256	0.160
Simple	2.385	0.418	0.692	0.676	0.242	1.504	0.388	0.661	0.619	0.214	0.461	0.256	0.256	0.163
spDistMult	0.620	0.218	0.701	0.582	0.202	0.679	0.243	0.759	0.568	0.189	0.448	0.423	0.249	0.181
spSimple	0.620	0.218	0.710	0.586	0.205	0.589	0.201	0.793	0.556	0.184	0.470	0.445	0.250	0.184




Figure 2. The percentage of reduction in epoch time achieved by spSimpleE compared to SimpleE for different negative ratios. The percentage is computed as $\frac{t_1 - t_2}{t_1}$ where t_1 is the time for SimpleE and t_2 is the time for spSimpleE.

of DistMult reduced by 2.7% and 18.4% on WN18AM and FB15k-237AM respectively whereas this reduction for spDistMult is only 0.4% and 0.8%. We got similar results for SimpleE and spSimpleE. These results show that the regularization term not only can act as a replacement for negative examples but also can have a regularization effect similar to that of dropout. Finally, we tested spDistMult and spSimpleE when computing Equation 3 for all entities and relations as opposed to only entities and relations in the batch. We observed only minor changes to filtered MRR (less than 1%) showing that computing the regularization term only

for entities and relations in the batch suffices.

Running times: As exemplified in Table 1, negative sampling usually consumes a large portion of the time for each epoch. Here, we compare the running time of SimpleE and spSimpleE for each epoch on the FB15k-237AM dataset. For a fair comparison, for spSimpleE we use a batch size $(n + 1)$ times larger than that of SimpleE as a rough estimate for memory usage (cf. Theorems 2 and 3). Figure 2 shows the improvement achieved by spSimpleE compared to SimpleE on FB15k-237AM in terms of epoch time. The improvement is measured as the percentage of reduction in epoch time. It can be observed that spSimpleE drastically reduces the epoch time especially for larger values of the negative ratio.

5. Discussion & Conclusion

In this paper, we proposed a novel regularization technique that obviates the need for negative examples during training and we showed its notable merits on several tasks and datasets. Future work can extend our regularization technique to other classes of models proposed in the literature, provide a deeper theoretical study of the proposed regularization and develop new regularization techniques, and extend the regularization to domains beyond relational embedding models where negative sampling is employed.

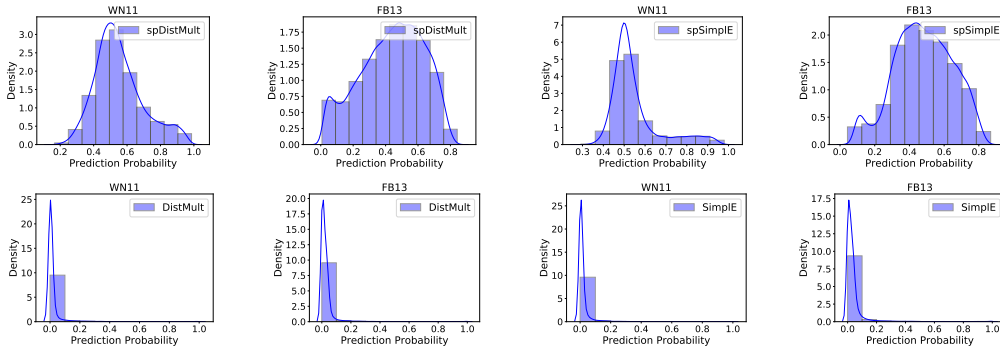


Figure 1. The probability outputs of different models on WN11 and FB13 grouped into bins as well as the kernel density estimates.

References

- Balkir, E., Naslidnyk, M., Palfrey, D., and Mittal, A. Using pairwise occurrence information to improve knowledge graph completion on large-scale datasets. *arXiv preprint arXiv:1910.11583*, 2019.
- Bollacker, K., Evans, C., Paritosh, P., Sturge, T., and Taylor, J. Freebase: a collaboratively created graph database for structuring human knowledge. In *ACM SIGMOD*. AcM, 2008.
- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. In *NeurIPS*, pp. 2787–2795, 2013.
- Brier, G. W. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950.
- Dettmers, T., Minervini, P., Stenetorp, P., and Riedel, S. Convolutional 2d knowledge graph embeddings. In *AAAI*, 2018.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 2011.
- Kadlec, Rudolf, O. B. and Kleindienst, J. Knowledge base completion: Baselines strike back. *arXiv preprint arXiv:1705.10744*, 2017.
- Kazemi, S. M. and Poole, D. Simple embedding for link prediction in knowledge graphs. In *NeurIPS*, pp. 4289–4300, 2018.
- Kazemi, S. M., Goel, R., Jain, K., Kobayez, I., Sethi, A., Forsyth, P., and Poupart, P. Representation learning for dynamic graphs: A survey. *JMLR*, 2020.
- Nguyen, D. Q. An overview of embedding models of entities and relationships for knowledge base completion. *arXiv preprint arXiv:1703.08098*, 2017.
- Nickel, M., Murphy, K., Tresp, V., and Gabrilovich, E. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2016.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- Platt, J. et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999.
- Ruffinelli, D., Broscheit, S., and Gemulla, R. You can teach an old dog new tricks! on training knowledge graph embeddings. In *ICLR*, 2020.
- Shan, Y., Bu, C., Liu, X., Ji, S., and Li, L. Confidence-aware negative sampling method for noisy knowledge graph embedding. In *2018 IEEE International Conference on Big Knowledge (ICBK)*, pp. 33–40. IEEE, 2018.
- Socher, R., Chen, D., Manning, C. D., and Ng, A. Reasoning with neural tensor networks for knowledge base completion. In *AAAI*, pp. 926–934, 2013.
- Tabacof, P. and Costabello, L. Probability calibration for knowledge graph embedding models. 2020.
- Terrell, G. R. and Scott, D. W. Variable kernel density estimation. *The Annals of Statistics*, pp. 1236–1265, 1992.
- Toutanova, K. and Chen, D. Observed versus latent features for knowledge base and text inference. In *Continuous Vector Space Models and their Compositionality*, pp. 57–66, 2015.
- Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., and Bouchard, G. Complex embeddings for simple link prediction. In *ICML*, pp. 2071–2080, 2016.
- Wang, P., Li, S., and Pan, R. Incorporating gan for negative sampling in knowledge representation learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Wang, Q., Mao, Z., Wang, B., and Guo, L. Knowledge graph embedding: A survey of approaches and applications. *IEEE TKDE*, 29(12):2724–2743, 2017.

Appendix

Proof of theorems:

Theorem 1. Let \mathcal{M} be a bilinear model with embedding dimension d where for each relation, only $d \leq \alpha \leq d^2$ of the elements in the embedding matrix can be non-zero. For a knowledge graph $\mathcal{G}_{\mathcal{E}, \mathcal{R}}$, the sum of scores of \mathcal{M} for all possible triples can be computed in $O(|\mathcal{E}|d + |\mathcal{R}|\alpha)$ as $(\sum_{e \in \mathcal{E}} \mathbf{z}_e^T)(\sum_{r \in \mathcal{R}} \mathbf{Z}_r)(\sum_{e \in \mathcal{E}} \mathbf{z}_e)$.

Proof. The sum of the scores of \mathcal{M} for all possible triples can be written as:

$$\begin{aligned}
 & \mathbf{z}_{e_1}^T \mathbf{Z}_{r_1} \mathbf{z}_{e_1} + \cdots + \mathbf{z}_{e_1}^T \mathbf{Z}_{r_1} \mathbf{z}_{e_{|\mathcal{E}|}} + \\
 & \mathbf{z}_{e_1}^T \mathbf{Z}_{r_2} \mathbf{z}_{e_1} + \cdots + \mathbf{z}_{e_1}^T \mathbf{Z}_{r_2} \mathbf{z}_{e_{|\mathcal{E}|}} + \\
 & \cdots + \\
 & \mathbf{z}_{e_1}^T \mathbf{Z}_{r_{|\mathcal{R}|}} \mathbf{z}_{e_1} + \cdots + \mathbf{z}_{e_1}^T \mathbf{Z}_{r_{|\mathcal{R}|}} \mathbf{z}_{e_{|\mathcal{E}|}} + \\
 & \cdots + \\
 & \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_1} \mathbf{z}_{e_1} + \cdots + \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_1} \mathbf{z}_{e_{|\mathcal{E}|}} + \\
 & \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_2} \mathbf{z}_{e_1} + \cdots + \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_2} \mathbf{z}_{e_{|\mathcal{E}|}} + \\
 & \cdots + \\
 & \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_{|\mathcal{R}|}} \mathbf{z}_{e_1} + \cdots + \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_{|\mathcal{R}|}} \mathbf{z}_{e_{|\mathcal{E}|}}
 \end{aligned}$$

Since matrix product is (left) distributive with respect to matrix addition, for each row of the above expression, we can factor the first two elements of the bilinear score and re-write the sum as:

$$\begin{aligned} & \mathbf{z}_{e_1}^T \mathbf{Z}_{r_1} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) + \mathbf{z}_{e_1}^T \mathbf{Z}_{r_2} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) + \cdots + \mathbf{z}_{e_1}^T \mathbf{Z}_{r_{|\mathcal{R}|}} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) + \\ & \cdots + \\ & \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_1} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) + \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_2} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) + \cdots + \mathbf{z}_{e_{|\mathcal{E}|}}^T \mathbf{Z}_{r_{|\mathcal{R}|}} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) \end{aligned}$$

Using the (right) distributivity of matrix product with respect to addition, for each row we can factor out the first vector of each element and re-write the above expression as:

$$\left(\sum_{e \in \mathcal{E}} \mathbf{z}_e^T \right) \left(\mathbf{Z}_{r_1} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) + \cdots + \mathbf{Z}_{r_{|\mathcal{R}|}} \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right) \right)$$

Using the (right) distributivity of matrix product one more time for the expression inside the right parenthesis, we can re-write the above expression as:

$$\left(\sum_{e \in \mathcal{E}} \mathbf{z}_e^T \right) \left(\sum_{r \in \mathcal{R}} \mathbf{Z}_r \right) \left(\sum_{e \in \mathcal{E}} \mathbf{z}_e \right)$$

Computing the expression above requires computing two sums one over all entity vectors and one over all relation matrices. Summing over all entity vectors has a time complexity of $O(|\mathcal{E}|d)$ and summing over all relation matrices has a time complexity of $O(|\mathcal{R}|\alpha)$ where α is the number of elements in relation matrices that are allowed to be non-zero. Therefore, the overall time complexity is $O(|\mathcal{E}|d + |\mathcal{R}|\alpha)$. \square

Theorem 2. *The time complexity of a single training epoch using Equation 1 in the main text is $O(\beta(1+n)|\mathcal{E}|\alpha)$.*

Proof. Assuming each entity appears on average in β triples, a single epoch requires computing model scores for $O(\beta|\mathcal{E}|)$ positive triples. With a negative ratio of n , a single epoch requires computing model scores for $O(\beta n|\mathcal{E}|)$ triples. Therefore, overall a single epoch requires $O(\beta(1+n)|\mathcal{E}|)$ score computations. Since the time complexity of computing a single model score is $O(\alpha)$, the overall time complexity becomes $O(\beta(1+n)|\mathcal{E}|\alpha)$. \square

Theorem 3. *The time complexity of a single training epoch using Equation 2 in the main text is $O(\beta|\mathcal{E}|\alpha)$.*

Proof. To prove this theorem, we first need to look at each mini-batch. For a mini-batch B , let \mathcal{E}_B and \mathcal{R}_B represent the entities and relations in the batch and β_B represent the average number of triples each entity in \mathcal{E}_B appears in. With a similar reasoning as in the proof of Theorem 2, the time complexity for computing the score for the positive triples in this batch is $O(\beta_B|\mathcal{E}_B|\alpha)$. The time complexity

for computing the regularization term is $O(|\mathcal{E}_B|d + |\mathcal{R}_B|\alpha)$. Note that $\beta_B \geq 1$ because every entity in $|\mathcal{E}_B|$ appears in at least one triple. We also assumed $|\mathcal{E}_B| > |\mathcal{R}_B|$ and $\alpha \geq d$. Therefore, $|\mathcal{E}_B|d < \beta_B|\mathcal{E}_B|\alpha$ and $|\mathcal{R}_B|\alpha < \beta_B|\mathcal{E}_B|\alpha$ and so the time complexity for computing the regularization term gets subsumed into the time complexity of computing the model scores and we can ignore it. By summing the time complexities for computing the model scores for all batches, we obtain the overall time complexity of $O(\beta|\mathcal{E}|\alpha)$. \square

Evaluation Metrics: To measure the performance on WN11 and FB13, following (Tabacof & Costabello, 2020), we report negative log-likelihood (NLL) and Brier score (Brier, 1950), two popular scores for computing how calibrated a binary classifier is. Formally, $NLL = \frac{1}{|Test|} \sum_{(h,r,t), l \in Test} \text{softplus}(-l * \phi(h, r, t))$ and $Brier = \frac{1}{|Test|} \sum_{(h,r,t), l \in Test} (\sigma(\phi(h, r, t)) - l)^2$ where the equation for NLL assumes $l \in \{-1, +1\}$ and the equation for Brier score assumes $l \in \{0, 1\}$. For completeness sake, we also report AUC.

To measure the performance on WN18AM and FB15k-237AM, for each triple (h, r, t) in the test set, we create two queries: $(?, r, t)$ and $(h, r, ?)$. For each query, we use the model to rank all entities in \mathcal{E} . Let $\kappa_{h,(?,r,t)}$ represent the ranking assigned to h for query $(?, r, t)$ and $\kappa_{t,(h,r,?)}$ represent the ranking assigned to t for query $(h, r, ?)$. We compute *mean reciprocal rank (MRR)* as the mean of the reciprocal of the ranks assigned to the correct entities and *hit@k* as the fraction of queries for which the correct entity is ranked among the top k entities. More formally, $MRR = \frac{1}{2*|Test|} \sum_{(h,r,t) \in Test} \left(\frac{1}{\kappa_{h,(?,r,t)}} + \frac{1}{\kappa_{t,(h,r,?)}} \right)$ and $Hit@k = \frac{1}{2*|Test|} \sum_{(h,r,t) \in Test} (\mathbb{1}_{\kappa_{h,(?,r,t)} \leq k} + \mathbb{1}_{\kappa_{t,(h,r,?) \leq k}})$ where $\mathbb{1}_{condition}$ is 1 if *condition* holds and 0 otherwise. The above setting is known as the *raw* setting. Bordes et al. (2013) advocated for a *filtered* setting where MRR and Hit@k are computed similarly except that the ranking $\kappa_{h,(?,r,t)}$ is only computed among those entities e where $(e, r, t) \notin Train \cup Validation \cup Test$ (and similarly for $\kappa_{t,(h,r,?)}$).

Implementation: To provide a fair comparison, we implemented our models and baselines in the same framework and measured their performance under fair experimental setup. The implementation is in PyTorch (Paszke et al., 2017) and AdaGrad (Duchi et al., 2011) is used as the optimizer. We selected the best hyperparameters using grid search with values selected as follows: learning rate $\in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$, λ (in Equation 2) $\in \{10^{-1}, 10^{-3}\}$, L2 regularization $\in \{10^{-1}, 10^{-2}, 10^{-3}\}$, embedding size $\in \{100, 200\}$, dropout probability $\in \{0.0, 0.4\}$, negative ratio $\in \{1, 10\}$, and ψ (in Equation 3) $\in \{-3, -2, -1, 0\}$. We fix I (introduced in the practical consideration paragraph of Section 3) to 5 and p (in Equa-

tion (3)) to 1. Note that the goal of our experiments is to show the merit of the proposed regularization compared to negative sampling as opposed to beating the state-of-the-art on a benchmark which requires much larger embedding dimensions, much larger grid for hyperparameter optimization, and adding several model-independent tricks (Kadlec & Kleindienst, 2017; Ruffinelli et al., 2020). For the experiments on WN11 and FB13, we selected the best hyperparameter setting and best epoch based on the NLL on the validation set. For the experiments on WN18AM and FB15k-237AM, we selected the best hyperparameter setting and best epoch based on the validation filtered MRR. The source code is available at <https://github.com/BorealisAI/StayPositive>.