

1 **All Reviewers:** Thanks for your insightful reviews. We have tried our best to clarify them below (common responses  
2 followed by individual responses). [A cryptic form of reviewer question/comment precedes our response.]  
3 [Experimental Results Interpretation] We confirm that Hits@1 and Hits@10 scores in Table 1 are correct. Note, it  
4 implies that our algorithm often ranks a ground truth entity either at Rank 1 or at a quite low rank. We believe, this is an  
5 artifact of our approach trying to assign different sub-spaces for different predicates and pushing the non-members away  
6 from them. This phenomenon also results in a poor score for (average) mean rank despite having high Hits@1 score.  
7 Such a behavior of E2R, as stated in paper, makes it practically interesting because Rank 1 candidate can be chosen  
8 with more confidence. [WN18 Dataset] Its a collection of binary relation instances where most of these relations satisfy  
9 transitivity property (e.g. hypernym) and have inverse relations (e.g. hypernym/hyponym). This dataset comes in the  
10 form of A-box and the underlying T-box capturing transitivity/inversion information is not supplied. Baseline methods,  
11 which are primarily distanced based, can probably capture transitivity/inversion properties from the A-box a bit better.  
12 [Deduction on LUBM] Deduction task is to deduce (pair of) entities that belong to a non-leaf (relation) concept in  
13 LUBM ontology. Each of 8 custom query is one such non-leaf (relation) concept – details at line no. 257-259 in paper.  
14 Training set include 69628 triples. Number of reasoning steps required can be as many as tree depth which is 4 here.  
15 **Reviewer #1** [Custom queries for LUBM] + [Results Interpretation] Please refer to line no. 3-14 above. [“T-Box” label  
16 in Figure 1] The label of the right figure in Figure 1 is actually correct. Both left and right figures contain “A-box” (oval  
17 nodes) as well as “T-box” (rectangle nodes) but the difference is that they depict unary and binary predicate hierarchy,  
18 respectively. We will clarify it. [Boldfacing Middle row in Table 1] As mentioned at line no 295–296 in paper, we  
19 chose to boldface only those entries which depict superior performance of E2R. We felt this may ease the comparison.  
20 [Abstract algebra terms] Yes, the field  $\mathbb{C}$  is space of complex numbers. We will clarify this as well as isomorphism.  
21 [Hilbert Space vs.  $R^d$ ]+ [Other work on inducing logical relation] Points well taken. We will address them.  
22 **Reviewer #2** [Expanding equations 4–9] We will include an explanation. [Operations in Section 3] These operations  
23 are trivial because we denote subspaces by indicator vectors and intersection, union, complement can simply be  
24 implemented via bit-wise AND, OR, FLIP operations. [Ontology Depth] For  $n$ -dimensional space, ideally, the ontology  
25 depth should be no more than  $n$ . [Higher-order Predicate] We also have a strong conviction that the trinary predicates  
26 can be represented by  $R^{3n}$ . However, other details (e.g. inner product definition) need to be worked out for this setup  
27 and that is our future direction. [Tuning of Baseline Parameters] Tuning of the parameters for the baseline approaches  
28 was performed on the dev set for FB15K and WN18 but on the training set for LUBM. For E2R, the tuning was always  
29 done on the training set. [HITS@1 and HITS@10 being same for E2R] Please refer to line no. 3–7 above. [Symbolic  
30 logic for link prediction] Yes, symbolic reasoners are typically not designed for link prediction tasks. [Symbolic  
31 reasoning vs. embedding] You are right in your assessment. We will surely include a discussion.  
32 **Reviewer #3** [Optimality of loss funct. (10)] We agree that locally optimal embedding given by SGD may not fully  
33 maintain the sanctity of the logical structure. However, because our formulation explicitly models sanctity, such a  
34 local-optima would still respect the sanctity to a greater degree (if not 100%) and hence would likely be better than  
35 that of the baseline approaches (as evident from LUBM results). Designing improved optimization technique is an  
36 interesting future direction. [Illustrative Example] We will try running through the example of Figure 2 for better  
37 explanation of eqn. (10). [Poincaré Embeddings + Lorentz Model] Thanks for pointing. By reading, we felt these  
38 methods do not make use of hierarchy information even if it is explicitly supplied in the form of T-box. Instead, they  
39 infer such hierarchical structure in an unsupervised manner. Therefore, these methods, by design, are not geared to  
40 tackle deductive reasoning problems where T-box plays a major role. We will include a discussion. [Comparison with  
41 32,33,34,35,10] [32] is less about embedding a KB and more about inducing an ontology. [33] embeds entities and  
42 predicates both in the forms of vectors which is very different from our philosophy. Unlike us, in [34], the training set  
43 upfront includes the assertions inferred on a select set of LUBM predicates (using symbolic reasoner) because their aim  
44 is to compete with symbolic reasoners (in speed and memory) during inference. [35] fall under the category of distance  
45 translation based embeddings (e.g. TransE) except that it uses ‘type information’ to improve quality of entities and  
46 relations embedding. Our Hits@10 score outperforms them on FB15K. In [10], they embed entities based on the type  
47 similarities as well as binary relations but they don’t explicitly preserve ontology of binary relations themselves. We will  
48 include these comparisons. [Sampling invalid entities vs. closed-world assumption] Like other approaches, we use a  
49 tiny fraction of invalid entities just to avoid degenerate solution, where all predicates (entities) collapse to a single space  
50 (vector). A full blown closed-world assumption involves many more negative assertions (e.g. missing edges among  
51 predicates) which we never touched. [“reasoning task” on LUBM?] + [Performance on WN18] Please see line no. 3–14  
52 above. [Finding all the entities] Yes, we enumerate all the entities. [Confidence/prob. corresponding to proj. length]+  
53 [“projecting each entity...non-fitment score”] Because E2R aspires to embed an entity within a concept subspace to  
54 which it belongs, we felt a natural loss metric would be residual length of its projection. The same idea is also endorsed  
55 by the Quantum Logic where projection length is indeed used as a probability of such membership assertion. We will  
56 clarify it. [better convergence with 3 -Ve entities] -Ve entities are defined at line no. 222-227 in paper and are used to  
57 improve convergence. Number of -Ve entities per +Ve entity is a hyper-parameter and value of 3 gave the best training  
58 accuracy. [Interpretability] Our formulation has separate loss for each aspect of preserving logical sanctity ( $\wedge$ ,  $\vee$ ,  $\neg$ ,  
59 entity membership, etc.). From these losses, one can argue where to tweak if sanctity is not maintained.