

1 Thank you to the reviewers for their detailed comments. We commit to addressing the minor and typographical errors.
2 In the following, we discuss the main issues raised:

3 **1. (R2) “if the d-DNNF is already compiled there must be a justification for why the logical embedding is needed”**
4 and “compare a more direct approach... like semantic loss (Xu et al 2018), and use the d-DNNF directly in the
5 loss to enforce constraints without embedding the formula.”

6 **Re:** As suggested, we compared our method against using d-DNNFs directly via the semantic loss proposed in (Xu et al. 2018)¹ on the visual relation prediction (VRP) task. Our method outperforms the semantic loss (Table 1).

Table 1: Comparison with semantic loss on visual relation prediction (VRP) task.

Model	Top-5 Acc.(%)
w/o semantic loss and w/o our embedder	84.30%
with semantic loss	84.76%
with GCN embeddings (ours)	92.77%

7

8 A potential reason for the performance difference is that the constraints in VRP are more complicated than those
9 explored in (Xu et al., 2018): there are thousands of propositions and a direct use of d-DNNFs causes the semantic
10 loss to rapidly approach ∞ . In contrast, our embedding approach avoids this issue. Moreover, our approach enables
11 generalization to constraints that involve previously unseen propositions; we can leverage representations such as
12 GLoVe word embeddings, which is not possible in the semantic loss. We will add these results and discussion to the
13 paper.

14 **2. (R2) “Building d-DNNFs is a difficult problem ... this difficulty is completely hidden in the paper”**

15 **Re:** Indeed, building d-DNNFs is a difficult problem in general and we will make this fact clearer in the paper. Practical
16 compilation has progressed significantly thanks to research in the area (e.g., the prior work pointed out by R2). We use
17 c2d (Darwiche, 2004) that can compile relatively large d-DNNFs in reasonable time; in our experiments, it took less
18 than 2 seconds to compile a d-DNNF from a CNF with 1000 clauses and 1000 propositions on a standard workstation.

19 We would like to emphasize that our GCN can embed other logic forms expressible as graphs. For cases where d-DNNF
20 compilation is not possible or prohibitive, other logic forms (e.g., CNF) could be used. The key contribution of our
21 approach is improved accuracy by embedding logical constraints. As suggested by the reviewer, we will highlight the
22 trade-off between accuracy and cost (to obtain compiled forms) in the final version.

23 **3. (R3) “Why not combine them jointly so that ... the framework becomes end-to-end?”**

24 Our framework can be trained end-to-end. However, we train the networks separately to (i) alleviate potential loss
25 fluctuations in joint optimization which makes training easier, and (ii) enable the same logical embeddings to be
26 used with different target networks (for different tasks). As suggested by the reviewer, the networks could be trained
27 end-to-end to fine-tune the embeddings further for a specific task.

28 **4. (R1&R3) “... more discussion of why the low and high cases don’t show as much improvement as the moderate
29 case”**

30 **Re:** We will provide additional discussion in the revision. In brief, we attempted to provide a good coverage of potential
31 problem difficulty. The “low” case represents easy problems that all the compared methods were expected to do well
32 on, and the “high” case represents very challenging problems that all methods were expected to struggle on. We posit
33 that the difference on the low and high regimes were smaller because (i) in the low case, all the methods performed
34 reasonably well, and (ii) on the high regime, embedding the constraints helps to a limited extent and points to avenues
35 for future work.

36 **5. (R3) “Comparison with some state-of-the-art neuro-symbolic methods, e.g. Logic Tensor Networks, could be
37 a plus”**

38 **Re:** Neuro-symbolic methods, such as Logic Tensor Network (LTN) and logic circuits, are not directly applicable in
39 our setting since the constraints for each input may differ. For example, in the visual relation task, the constraints for
40 each image was different. One possible approach would be to conjunct all the constraints, but there would be 7×10^5
41 propositions and $\approx 10^5$ clauses, which results in a prohibitively large logic network.

42 **6. (R3) “Provide more details on the methodology such that it could be better understood and enhance its
43 reproducibility.” and “explain how they chose some of their parameters.”**

44 **Re:** We used grid search to find reasonable parameters; we will state this and provide additional methodological details.
45 We have also submitted our implementation code which can be used to reproduce all the results in the paper.

¹We used publicly available code provided by the authors.