To reviewer #1:

Broader applicability of the method. Our method can be applied to a broader range of tasks that fall into visual reasoning. To name a few, Newtonian physics problem solving task [1], geometric problem solving task [6] and CLEVR as reviewer #3 suggested. We have observed a growing number of reasoning tasks in NLP such as multi-hop text question answering, knowledge graph reasoning and conversational models. Our approach can potentially be applied to strengthen the reasoning abilities of such tasks.

Availability of the problem categories. Category is widely available for abstract reasoning / visual reasoning which state-of-the-art models leverages as type loss. When such categories are missing, one can either get it using an unsupervised way such as the paper[5] that reviewer #2 suggested. Another way to deal with such a problem is to assume a latent category variable and optimizes it together with teacher model which could be a promising future work.

Differences to [7]. We compared methods used in [7] and as shown in table 1 and table 2, it is not performing well for abstract reasoning. In [7], only the loss of the last step is used as opposed to the complete trajectory in our method. Additionally, in [7] action is taken to be a 0/1 decision on sample id but our action is a proportion of the related problems. All of these leads to the better performance of our teacher model. We will improve the paper to clearly outline the differences.

Confirmation of the source of performance gain. We have illustrated in the empirical study that training with a specific trajectory with different proportion of the distracting/reasoning features can dramatically improve model performance. This is illustrated in Figure 1, which shows that the difficulty of abstract reasoning task lies in the existence of distracting features. Table 1 shows that the appropriate training trajectory can greatly improve the model in the presence of distracting features. At the same time, the visualization of Sec 5.4 also shows that our method can distinguish distracting features better. We will put a visual training trajectory map in the final paper for better illustration.

Other issues. The reviewer is right that χ_k (L151) is a set of embedding of all triple-panels. We will fix it.

To reviewer #2:

Disentangled representations. Disentangled representation separates information on a single input while our method select inputs for a model. Our method is orthogonal to disentangled representation and can be applied on top of it. It is also worth mentioning that [4] actually implemented the idea of disentangled representation but with little improvements on abstract reasoning. One intuitive explanation is that distracting features live in a much more illusive manifold and disentangled representation along is not capable of separating it from reasoning features.

Related work. We will add discussions of disentangled representation into our related work.

Extrapolation experiments. We actually did an extrapolation experiment on the PGM. We separated training and testing in a way that they have non-overlapping values of "color" attribute and have achieved 8% improvement in accuracy. We are happy to include this result along with additional experiments in the final paper.

Performances of models other than LEN. We actually have a complete set of comparisons of WReN with/without teacher model in table 2. Performance of WReN is improved from 75.6% to 77.8% with type loss and from 62.8% to 68.9% without type loss. We will complete comparisons of other baselines (e.g., RN) in the final paper.

To reviewer #3:

Performance improvements of the teacher network on LEN compares to WReN model on PGM. The performance improvements on WReN is as significant as the one on LEN. Since the codebase of WReN has never been released, we implemented it on our own with an accuracy of 70.1% using type loss but we have never been able to reproduce the reported benchmark (i.e., 75.6%). Nevertheless, we include the performance of the published results of WReN in our paper for fair comparison. As a matter of fact, if we compare the improvements of teacher model against our own baseline on WReN the improvement is 7.2%. Comparing to 11.1% of improvements with LEN model.

Testing on additional visual reasoning tasks. We actually did an experiment on the CLEVR dataset but we didn't include it into the paper. our LEN model achieves 1.7% accuracy increase(95.5% to 97.2%) compared to RN[2]. Please note that CLEVR dataset is much easier than the datasets we used in the paper and the already high performance on baseline method allow only a marginal improvements using our teacher model. Nevertheless, the performance gain seems to be significant an consistent. We will explore the efficacy of the model on more widely used visual reasoning tasks (E.g., CLEVR-CoGeNT task) in the final paper.

Writing Issues. We will fix these issues as the reviewer suggested.

[1] Sachan M, et al. Parsing to programs: A framework for situated qa. KDD 2018.
[2] Santoro A, et al. A simple neural network module for relational reasoning. NIPS 2017.
[4] Steenbrugge, X., et al. Improving generalization for abstract reasoning tasks using disentangled feature representations. In Workshop on NIPS, 2018.
[5] Hsu, Kyle, et al. Unsupervised learning via meta-learning. arXiv 2018.
[6] Seo M, et al. Solving geometry problems: Combining text and diagram interpretation. EMNLP 2015.
[7] Yang Fan, et al. Learning to teach. ICLR 2018.