We thank the reviewers for the reviews. We are glad to see that the reviewers liked the significant and novel contributions of the paper. We address specific questions and misunderstandings in the reviews.

Reviewer #1

On train-test split ratio: We refer the reviewer to ***label rate*** on Table 3 in our paper. These numbers are precisely the $\frac{train}{train+test}$ split ratios on each dataset.

On a robust experimental setting: As mentioned in the paper, we have reported mean accuracy and standard deviation over 100 different train-test splits in semi-supervised learning experiments to ensure a robust experimental setting. We refer the reviewer to the supplementary material/appendix for more experiments on different train-test split ratios. Moreover, as confirmed by Reviewer $\#2$, extensive experiments in both semi-supervised learning and combinatorial optimisation demonstrate the effectiveness of our methods.

On comparison against GCN: We have thoroughly compared against HGNN [3] in all our experiments. HGNN uses the clique expansion [4]. Hence, it is exactly the GCN baseline that is easily extended to hypergraphs [1].

On typos: Thanks for pointing the typos out. We will fix them in the final version.

Reviewer #2

On the lack of clarity: Many thanks to the reviewer for pointing out the lack of clarity in the paper. We will include background reviews of all the new hypergraph Laplacian operators in the final version of the paper by providing a crisp overview for each work. We will ensure that the paper is self-contained so that the broader NeurIPS community can appreciate the significance of our work better.

On 1-HyperGCN vs. FastHyperGCN: It is to be noted that 1-HyperGCN samples one edge (in each epoch) while FastHyperGCN does use the mediators and hence samples $2|e| - 3$ edges for each hyperedge $e \in E$. Hence, it is quite intuitive that FastHyperGCN is superior to 1-HyperGCN.

On different weights on edges: We observed that uniform weights on all edges do give the best results compared to other weights (e.g. zero weight for $\{i_e, j_e\}$ and uniform weights for the remaining edges to the mediators). As suggested by the reviewer, we will include a comparison table in the appendix.

On small issues: Thanks for pointing these out. We will fix them in the final version.

Reviewer #3

On insights into the proposed method: We refer the reviewer to Section 6 of our paper for insights into the method. As confirmed by Reviewer #1, we have provided theoretical analyses on the results for insights.

On when and why HyperGCN outperforms HGNN: We refer the reviewer to Table 5 in our paper which provides insights into when HyperGCN outperforms HGNN [3]. Our methods produce sparser approximations which accumulate less noise and hence are superior on nosiy datasets as shown in Table 5.

On comparison against more recent hypergraph methods: We have thoroughly compared against the two known state-of-the-art semi-supervised learning methods on undirected hypergraphs viz., HGNN [3] and explicit Laplacian regularisation [2] (MLP + HLR in our paper).

References

[1] Sameer Agarwal, Kristin Branson, and Serge Belongie. Higher order learning with graphs. In *ICML*, 2006.

[2] T.-H. Hubert Chan and Zhibin Liang. Generalizing the hypergraph laplacian via a diffusion process with mediators. In *COCOON*, 2018.

[3] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. In *AAAI*, 2019.

[4] Denny Zhou, Jiayuan Huang, and Bernhard Schölkopf. Learning with hypergraphs: Clustering, classification, and embedding. In *NIPS*, 2007.