

1 We thank the reviewers for detailed and constructive reviews and respond to each reviewer’s questions:

2 **R1: Utility of Strong Generalization:** Generalization and strong generalization on out-of-distribution data (co-
3 variate shift) are fundamentally unsolved problems in both machine learning and specifically program synthesis/neural
4 execution. Our insight that faithfully learned feedback masks helps strong generalization of transformers in a difficult
5 domain is an important contribution to understanding why strong generalization occurs.

6 **Importance of Strong Supervision:** Our primary objective with this paper is to establish whether strong generalization
7 can be achieved (to a reasonable degree) in a mechanistic sense. This is effectively a form of imitation learning. Followup
8 work would focus on relaxing this constraint, but we believe that this is beyond the scope of a single paper. Other papers
9 in the field have taken a similar path, for example [21] explores program synthesis via strong supervision and Pierrot et
10 al., 2019 build on this with an RL solution called AlphaNPI that uses weak supervision.

11 **[29] Comparison:** Code for [29] is not available, making faithful reproduction difficult. NEE & [29] are based
12 on different mechanisms. At each step, [29] requires supervision on every node and edge, while NEE instead
13 supervises on the mask and nodes. We tried to match their setting (using the same graph types) and define our
14 accuracy using a stricter metric (an exactly correct whole-graph result instead of next-node accuracy). For MST,
15 [29]’s next-node accuracy at 100 node graphs is 41.37%, implying lower accuracy with the strict whole-graph correct
16 metric vs close to 100% for NEE. There are evaluation differences ([29] trains using 20-node graphs, NEE trains on 8-node graphs), but we feel our evaluation vs. [29] is
17 fair. [29] suffers from attention degradation which they try to counteract with max aggregation and entropy penalties.

24 **Binary Ablation:** We provide a binary ablation in the Appendix: (Fig.7, all_modifications variant V.S.
25 all_modifications_one_hot_emb variant). **Figure-5 Insight:** We don’t believe that it’s expected for close numbers to
26 necessarily show structure due to PCA. Figure-R 1 shows the embedding projections before and after training.

27 **R2: Motivation:** Briefly, we establish that, even with strong supervision, transformers can’t generalize on simple
28 algorithmic tasks. We propose a learned conditional masking mechanism as a fundamental technique to improve this.
29 **Results:** We can’t say that performance will be perfect in general, but we observed this behavior in our experiments.
30 On the MST task, we do notice some degradation over large test graphs, but the degradation is graceful. **Figure-5:**
31 We’ll work on numerical analysis. Our goal was to establish that a) these embeddings have structure, b) these structures
32 differ meaningfully across tasks, and c) non-observed numbers are embedded correctly into their appropriate spots.

33 **R3: Baseline Transformer Masks:** We apologize for not being clear here. Yes, both NEE and the baseline transformer
34 have the same strong supervision granularity (output value & correct mask). During training, the baseline transformer
35 is also given the ground truth masks (Figure 3). During testing, masks are not provided for both models and the
36 baseline’s attention mask consequently generalizes poorly (Figure 4). NEE’s improvement is due to the learned mask
37 and architectural design. A plain transformer without mask modifications at both training and testing times performs
38 poorly (which is why we started working on NEE). We’ll include this ablation in the next version of the paper.

39 **Figure 4:** The element in the (x, y) grid indicates the strength of the attention score between the x th number and the y th number. Figure-4 shows that the attention in the
40 first 8 rows of the baseline transformer (below training data length, attention between the smallest 8 numbers) is clear, but out of distribution attention (beyond 8 numbers)
41 grows fuzzy, resulting in incorrect predictions.
42
43

44 **Performance:** Yes, we find that using subroutines that strongly generalize, the resultant composition also strongly generalizes. In Table-1 we evaluate a challenging setting
45 (training on 8-nodes, testing on 100). In two cases for shortest path, accuracy isn’t exactly 100%, but 99.99% and 99.98%, which we don’t consider significant deviations.
46 In Table-5, using the learned subroutines we achieve 100% accuracy - scenario 1 (S1, first two table rows). You refer to the second two rows (S2), where we train using only Erdős-Rényi random graphs but
47 test with a mix of 4 kinds of graphs (Section A.3, Appendix). In this case, the test distribution drifts from the training distribution (Figure-R. 2), and performance suffers, although the resultant accuracy is still much higher than [29].
48
49
50
51

52 **Holding Out Numbers:** Yes, the numbers are randomly selected. You are correct that if a bit is not toggled (i.e. bit 6
53 for 64 in your example), the model would not be able to generalize, but this motivates future work to create numerical
54 basis that extrapolate across unseen bits, whereas our focus here is on interpolation.

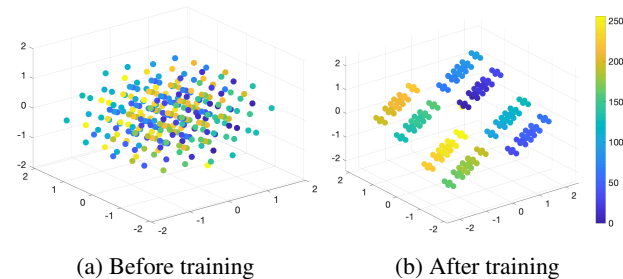


Figure-R 1: PCA projections on number embeddings




Figure-R 2: Histograms of data used in MST (Setting 2)