arXiv:2107.12580v1 [cs.LG] 27 Jul 2021

Pointer Value Retrieval: A new benchmark for understanding the limits of neural network generalization

Chiyuan Zhang[*1] **Maithra Raghu**[*1] **Jon Kleinberg**[2] **Samy Bengio**[†1]

[1]*Google Research, Brain Team*
[2]*Cornell University*

Abstract

The successes of deep learning critically rely on the ability of neural networks to output meaningful predictions on unseen data — *generalization*. Yet despite its criticality, there remain fundamental open questions on how neural networks generalize. How much do neural networks rely on *memorization* — seeing highly similar training examples — and how much are they capable of human-intelligence styled *reasoning* — identifying abstract rules underlying the data? In this paper we introduce a novel benchmark, Pointer Value Retrieval (PVR) tasks, that explore the limits of neural network generalization. While PVR tasks can consist of visual as well as symbolic inputs, each with varying levels of difficulty, they all have a simple underlying rule. One part of the PVR task input acts as a *pointer*, giving the location of a different part of the input, which forms the *value* (and output). We demonstrate that this task structure provides a rich testbed for understanding generalization, with our empirical study showing large variations in neural network performance based on dataset size, task complexity and model architecture. The interaction of position, values and the pointer rule also allow the development of nuanced tests of generalization, by introducing distribution shift and increasing functional complexity. These reveal both subtle failures and surprising successes, suggesting many promising directions of exploration on this benchmark.

1 Introduction

The remarkable capabilities of neural networks across many different domains is critically reliant on their ability to *generalize* — output accurate predictions on *unseen data*. But despite its central importance, the mechanisms of neural network generalization are still poorly understood, with many unanswered questions. For example, in [35], it is shown that neural networks can surprisingly *memorize* training data with random labels. Is this a barrier to true, intelligent, generalization? Followup work paints a more nuanced picture, showing that like in humans, memorization can also aid generalization [15, 9]. But humans are also perform generalization through alternate approaches, such as abstract *reasoning*. Are general neural networks also capable of this?

Answering these questions on neural network generalization has gained pressing urgency as we increase the complexity of tasks being considered (e.g. object detection for self-driving, multimodal learning) [11, 8, 24], run into limits on the amount of diverse data available [3], and also require reasoning capabilities in high-stakes settings, e.g. medical data with distribution shifts [33, 21]. But progress in understanding generalization has been hindered by a lack of benchmarks that can distinguish between simple methods of generalization such as memorization or nearest neighbors from more complex approaches such as sophisticated reasoning.

---

[*] Equal Contribution
[†] Currently at Apple

In this work, we address this shortcoming, introducing a simple but versatile benchmark to test and understand the limits of neural network generalization. Specifically, our contributions are:

- We introduce Pointer Value Retrieval (PVR), a family of tasks to study neural network generalization. PVR tasks can consist of visual or vectorized inputs, with natural ways to systematically vary task difficulty by introducing distribution shifts and/or increasing functional complexity.
- All PVR tasks are unified by a simple rule, where part of the input acts as *pointer*, indicating a different location in the input which forms the *value* (output). This property means that PVR tasks require understanding of position, values and the rule connecting them, which naturally enables many rich task variations. It also allows targeted testing of which of these abstract concepts a model has learned, and which it has failed at.
- We perform an empirical exploration of neural network generalization, starting with visual PVR tasks and introducing distribution shift. We analyze ways in which this causes models to fail, identifying important questions for progress on these tasks.
- We then look at neural network generalization as the functional complexity of PVR tasks are varied. We introduce a family of functions to aggregate values into final output, and highlight connections between this and the notion of noise sensitivity from Boolean functions, which measure task difficulty.
- We test multiple model architectures on vectorized PVR tasks with varying functional complexity, discovering performance differences between MLPs, Transformers and the recent MLP-Mixer.
- We analyze the models for abstract concepts learned, discovering surprising successes on tasks combining functional complexity and distribution shift. We discuss the ramifications and important open questions for future exploration.

2 Related Work

Neural network generalization and reasoning has been studied with several approaches/tasks distinct from our PVR framework. One active direction has looked at the capability of models to learn visual tasks in a few-shot setting [29, 31], on benchmarks such as Omniglot [18], and Bongard-LOGO [19]. Our focus is not on few-shot, but faithfully learning the reasoning rule. Previous work has looked at learning rules in visual question answering [1, 13], drawing on classical work studying generalization [10]. These rules are often much more complex than our simple pointer-value formulation. Other work with varied rules has looked at mathematical reasoning, both in the form of non-verbal, visual puzzles [2, 36] as well as more symbolic mathematical inputs [28, 27], with related efforts studied in RL [25, 35, 23]. There are also tasks exploring verbal reasoning abilities [17, 20], which are different from our setting with systematic difficulty variation. Sequence-to-sequence learning has been used to study compositional generalization [17, 14, 26], due to rapid advances in natural language processing in the past few years. However, the model architectures and learning dynamics are much more complicated to analyze than the classification setting. A line of literature design neural networks with dedicated symbolic components to improve systematic generalization [4, 5]. We instead focus on studying general architectures and analyze whether generalization behavior is learned from the data.

3 Pointer Value Retrieval

Here we introduce a new benchmark, *Pointer Value Retrieval (PVR)* for testing and understanding neural network generalization. The PVR benchmark provides a rich family of tasks with different inputs types and varying difficulty. But all of these tasks are unified by a simple underlying rule: part of the input acts as a *pointer*, indicating a position in the input, which contains the *value*, which forms the output. This structure requires models to understand notions of position, values and the rule linking them, enabling a diverse set of task types. We overview different subfamilies of tasks below.

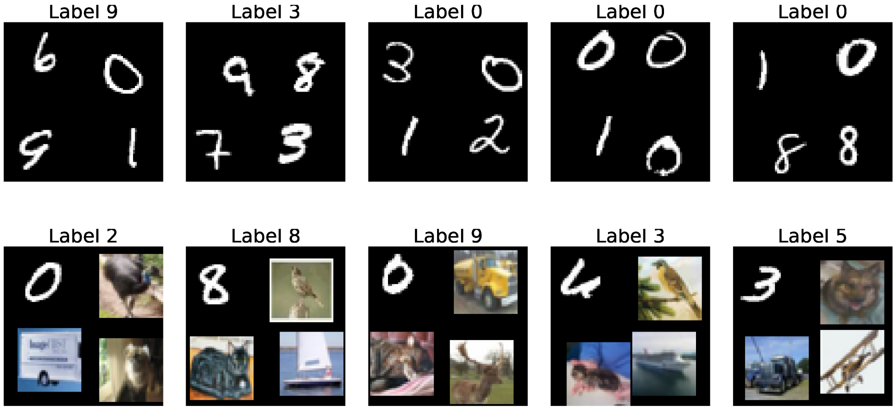

Figure 1: **Block style visual PVR tasks.** In the above examples, we have four images arranged in a 2x2 grid. The top left image acts as a pointer, with the actual digit indicating which of the remaining images is the value. In the examples above, we take pointer digits in $0 - 3$ to indicate the upper right position, $4 - 6$ the lower left position, and $7 - 9$ the bottom right position. The image at that position is then the value, which determines the label (the digit value for top row of MNIST examples, and class label for bottom row of CIFAR-10 examples.) Task variants may have larger block sizes (e.g. 3x3 blocks) or a different mapping from pointer digits to positions. Train, validation and test splits are easily created by drawing images from the underlying dataset train/test splits.

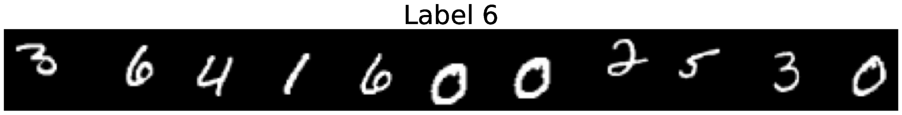

Figure 2: **Sequential style visual PVR tasks.** Visual PVR tasks can also be sequence style, with one token acting as a pointer, and determining the position of the value. In the above examples, the left most digit is taken as the pointer, with the digit simply indicating which of the remaining 10 positions (indexed with $0 - 9$) will be the value, from which the label is derived. Other variations may change the length of the sequence, vary the location of the pointer digit, or add in different datasets.

3.1 Visual Pointer Value Retrieval

One subfamily of PVR tasks, visual PVR, consist of visual (image) inputs, which provide a way to test and understand generalization in high performing computer vision models. We consider two structures for our PVR visual inputs (i) *block* and (ii) *sequential*, illustrated in Figures 1, 2.

In the block styled input, we have a 2x2 grid, where the upper left entry acts as the pointer. The number of the pointer indicates which position of the input we should look at for the *value* term, the label of which provides the overall label. For example, in the top row of Figure 1, we define that if the pointer digit is $0 - 3$, the value term is the upper right digit, if the pointer digit is $4 - 6$, the value is the lower left digit, and if the pointer digit is $7 - 9$ the value is the lower right digit. This task structure also extends across datasets (shown with MNIST digits as pointers and CIFAR10 images as possible values in Figure 1 second row), and to a larger block size.

In the sequential input, we have the visual data arranged in a row. In our example, we take the left most entry acting as the pointer, and the numerical digit specifying which of the remaining sequence will be the value. In Figure 2, we show this for a sequence of length 11 (one pointer and ten possible values), but like the block styled input, this can also be extended to different sequence lengths, pointer positions and different datasets.

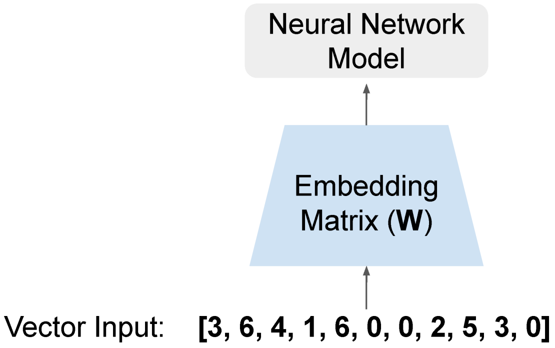

Figure 3: **Using an embedding layer with vectorized PVR tasks.** PVR tasks can also take in vector inputs, to disentangle considerations of generalization from visual representation learning. In our settings, these vector inputs consist of digits, similar to the visual tasks. There is an analogy to language settings, with the range of digit values forming a vocabulary, and indeed, when developing models for vectorized PVR tasks, we first embed all the digits with an embedding matrix.

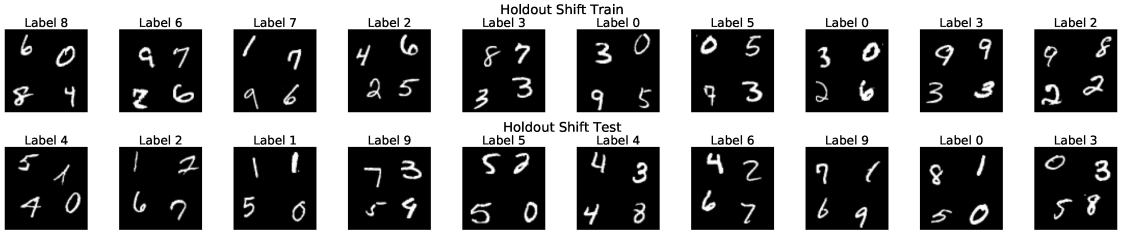

Figure 4: **Introducing distribution shift into block style PVR task.** We introduce distribution shift by holding out some digits at some positions during training. Specifically in the train set, digits $1-3$ never appear in the top right, $4-6$ never appear in bottom left, $7-9, 0$ never appear in bottom right (top row). At test time, we can test on examples drawing on all $0-9$ iid in all positions, or on the *holdout shift*, exactly the examples that were held out at training (bottom row).

3.2 Vectorized Pointer Value Retrieval

To disentangle learning visual representations from learning reasoning, we can also look at *vectorized* versions of these tasks, where the input is no longer an image, but a *vector* consisting of the digits themselves. This has natural parallels to language tasks — the digits form a vocabulary. Indeed, when testing models on this benchmark, we also use an embedding layer, which maps each digit to an embedding vector, as show in Figure 3. As with the visual input, in our examples we let the left most entry act as the pointer, indicating which of the remaining digits is the value (and output).

3.3 Varying Difficulty with Distribution Shift

Having seen the different input types and the basic tasks, we turn to defining how to systematically vary the difficulty of the PVR tasks. It is this feature that makes the PVR tasks especially useful in testing and understanding generalization. Specifically, as the PVR tasks rely on reasoning about position *and* values, there are several ways to naturally increase task complexity.

One such method is to introduce distribution shift: at train time, we (i) ensure some of the values do not appear at some positions (i.e. value *holdout* at some positions) (ii) we also ensure all values appear somewhere (so the model has the ability to learn all the values). We can then test to see how well the model learns notions of position, value and the pointer rule. In Figure 4, we provide an example of this (studied further in Section 4): we take the block style visual PVR task, and in the training set, make sure digits $1-3$ never appear in the top right, digits $4-6$ never appear in the bottom left and digits $7-9, 0$ never appear in the bottom right.

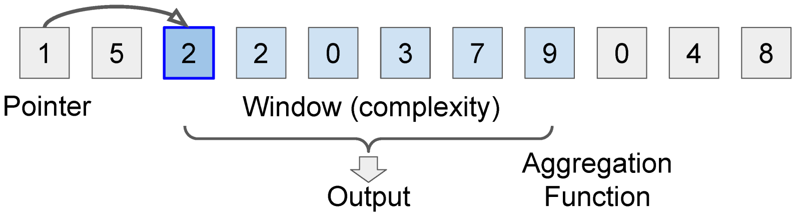

Figure 5: **Increasing functional complexity of PVR tasks.** Instead of the pointer indicating a single position which becomes the value (and output), the pointer instead indicates a *window* of values, which are aggregated to form the final output. The size of the window can be determined by a complexity parameter. The schematic has window size 6.

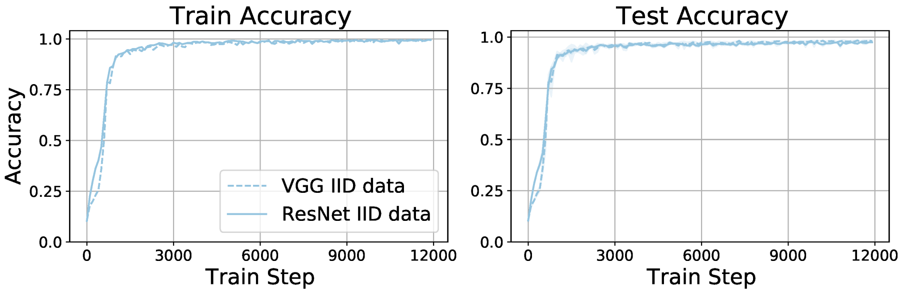

Figure 6: **Neural networks successfully generalize on the block style Visual PVR task.** We train two standard convolutional networks, VGG and ResNet, on the block style Visual PVR task depicted in Figure 1 (row 1). We find that both models learn to generalize very well (right pane).

3.4 Varying Difficulty with Functional Complexity

Another natural method to vary task difficulty is by increasing *functional complexity*. So far, the pointer in the input has indicated the position of the value, which has *directly* been taken as the output. Instead, we now use the pointer to indicate a *window* of values, which are aggregated to form the final output. The size of the window is determined by a fixed, complexity parameter, with window size 0 corresponding to the original task of directly outputting the value. Figure 5 shows a schematic of this.

The connection between window size and functional complexity can be made formal through Boolean function analysis (discussed in Section 5.1). But it can also be seen through observing that having a window of values increases the number of positions used to determine the output. This prevents reliance on naive nearest neighbors, as a single digit changing could affect the output.

There are many possibilities for aggregation functions when increasing functional complexity. In our experiments, we study the following, discovering interesting variations in difficulty amongst them: (i) ***mod_sum*** (sum all the values in the window and compute the remainder mod 10) (ii) ***median*** (median of window values) (iii) ***maj-vote*** (mode of window values) (iv) ***min/max*** of window values.

4 Warmup: Initial Exploration of Visual PVR and Distribution Shift

Having introduced the PVR benchmark, different subfamilies of tasks it contains, and natural ways to vary difficulty, we turn to a simple first empirical study of the block style visual PVR task, and the effect of introducing distribution shift. In Figure 6, we look at training two convolutional networks, VGG and ResNet, on the standard block style visual PVR MNIST task. (Full details of all PVR datasets in Appendix Section A.) We observe that both model architectures achieve strong generalization (Figure 6 right pane.) But how is this

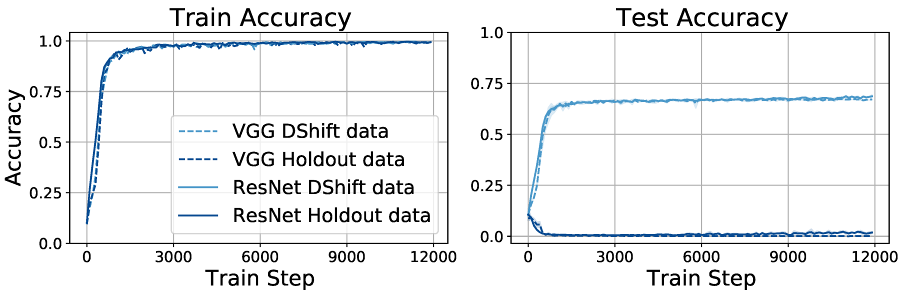

Figure 7: **Introducing distribution shift into block style PVR results in generalization failures.** We introduce distribution shift by witholding some digits from some positions during the training set, as illustrated in Figure 4. The network still sees all pointer digits and all digits in some positions, so has enough information to learn notions of digits, positions and the pointer rules. However, evaluating on (i) *Dshift* a test set where all digits again appear in all positions (ii) *Holdout* a test set of precisely all examples with digits in locations holdout during training (pictured in Figure 4 bottom row), shows failure to generalize. Particularly in Holdout, we see the model converges to 0% accuracy from random chance at start.

performance achieved? Are the models learning the notion of the four different positions, the ten different digit values and the pointer rule?

To understand this further, we introduce distribution shift into the task, specifically the holdout rule described in Figure 4, where some digits constitute the holdout in some positions at training, but all digits (and all pointer values) appear. This provides the models with enough information to still learn the abstract concepts (position, digit values, pointer rule) underlying the labelling process, but prevents simple memorization or nearest neighbor approaches.

We then evaluate the models on data which draws uniformly at random from $0 - 9$ at all positions (labelled *dshift*) and also all the datapoints that were heldout during training (labelled *holdout*.) The results, shown in Figure 7, are striking. We see that both models perform poorly when tested on the *dshift* data, and entirely fail to learn (converging from random accuracy to 0%) when tested only on the holdout data. This reveals that strong generalization performance seen in Figure 6 is not due to the neural network correctly learning to reason, but a consequence of a more simplistic generalization method, like nearest neighbors.

The striking convergence to 0% accuracy in the Holdout case also suggests a systematic error. To understand this further, we analyzed the model representations and outputs. In particular, looking at the raw logit values across models trained in an IID setting and in the Holdout setting revealed a surprising insight, shown in Figure 8. We observe that under Holdout shift, models learn to negatively correlate the pointer digit with values that were held out during training. In Figure 8 left pane, as digits $1 - 3$ never appear in the top right position during training, and pointer 0 indicates the top right position, the neural network automatically assigns labels $1 - 3$ a very low logit value, even though at test time, all of the top right values are in $1 - 3$. By contrast, in the IID setting (right column) no such adversarial correlation is observed. Additional examples are in the Appendix.

These results raise interesting open questions that can be explored on this benchmark. Are there architectural changes that can enforce better priors and withstand distribution shift? Can novel learning objectives prevent these adversarial correlations? Progress on these questions holds promise for greater robustness.

The results also highlight a question on natural mechanisms of neural network generalization. On the one hand, the neural network did not learn to generalize on this task with the abstractions we might have expected, and new methods that better align its generalization to these concepts could be very useful. However, its systematic failure arose from spotting a pattern relating pointers to labels that was actually present in all of training data. This suggests looking at settings of increased functional complexity, where the additional steps in mapping value to label remove such patterns.

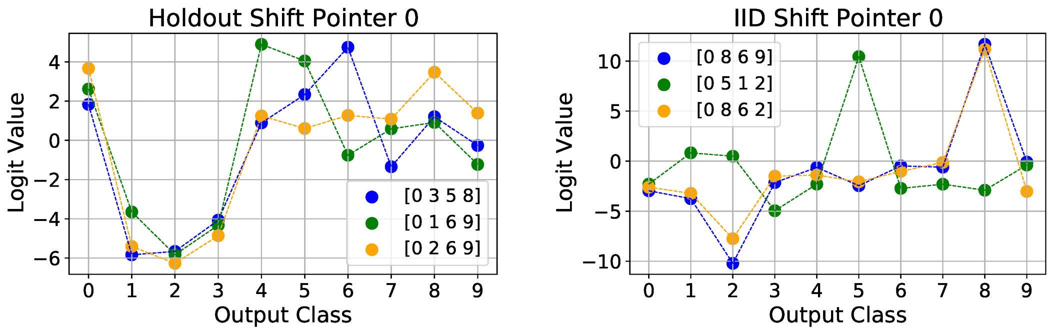

Figure 8: **Analyzing the failure of neural networks on Holdout shift reveals systematic mistakes from learned correlations between pointer digits and labels.** Left: we plot the raw logit values for a few different test examples for pointer digit 0. We observe that the model has learned to assign very low logits to labels $1 - 3$, exactly the values left out from the top right position during training (which pointer 0 points to). Although *all* test examples have only values $1 - 3$ in this position, this correlation is ingrained in the network, leading to systematic errors. For comparison, we include logits from models trained in the IID setting, where we observe no correlations between pointer digit and label values. Additional examples are in the Appendix.

5 Functional Complexity and Noise Sensitivity

Motivated by the previous results, we turn to studying performance and generalization on PVR tasks of increasing functional complexity. Are there limits to the functional complexities neural networks can generalize to? And how does dataset size influence this? Are there inherent difficulties in different aggregation functions, and can we formalize this? We investigate these questions below.

Our empirical investigation centers on PVR tasks with vectorized inputs (Section 3.2), where functional complexity is varied through increasing window (neighborhood) size, with neighborhood size 0 corresponding to the original setting of the pointer indicating a value that directly becomes the output. We start using an `MLP` architecture, where the terms in the vector are first embedded with a learned embedding matrix (Figure 3), and use the *mod_sum* aggregation function due to it being most sensitive to the neighborhood values (a single change of any value changes the output.)

Figure 9 illustrates the results, with some striking takeaways. Firstly, we note that training accuracy across all complexities (neighborhood sizes) and dataset sizes is consistently 100%. This contrasts with test accuracy which starts off random at lower training set sizes and shows a thresholding behavior, jumping up to 100% with large enough training data. (Neighborhood 3 is an exception, where even at 10^6 training points, test accuracy remains at random.) Together, this reveals that at small dataset sizes, these models completely overfit, and are not learning the underlying reasoning rules at training time. With enough data however, we see their test accuracy jumps up. In Section 7, we dive further into this, exploring whether 100% test accuracy translates to learning the right reasoning rules.

5.1 Measuring Complexity with Noise Sensitivity

We apply *noise sensitivity*, used to determine complexity of boolean functions [22, Proposition 3.3], to quantify the complexity of our tasks. Intuitively, the noise sensitivity $\mathrm{NS}_\delta[f]$ measures how sensitive the outcome of f is to random perturbations with probability δ of the input bits. Figure 10a plots $\mathrm{NS}_\delta[f]$ with a range of δs, and the results are consistent with our intuition that target functions with higher number of neighbors are more complex. In fact, the choice of the *mod_sum* label aggregation is also important. As shown in Figure 10b, other common aggregations such as majority voting don't generate a clean sequence of tasks with increasing difficulties. For `min` and `max` aggregation, the tasks actually become easier as neighborhood size increases.

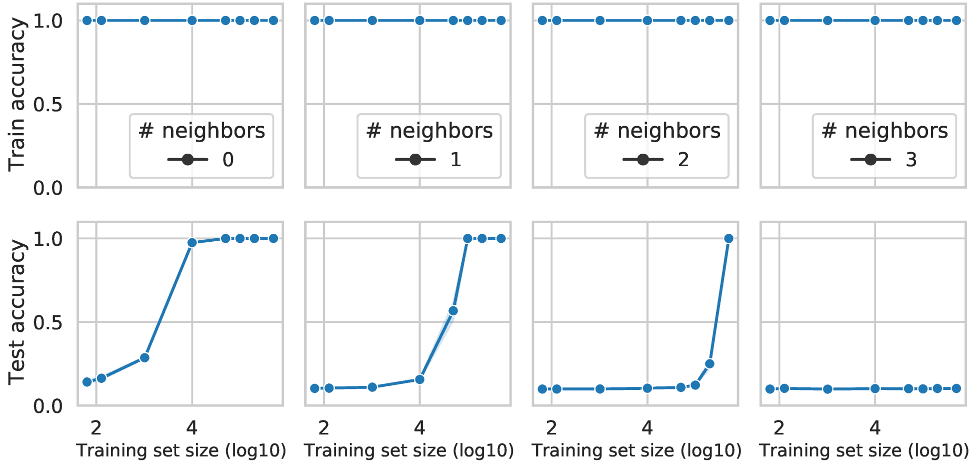

Figure 9: **Performance evaluation on PVR tasks when increasing functional complexity and dataset size.** We train `MLP` models on vector input PVR tasks of varying complexity (neighborhood size) and training set sizes. While training accuracy is consistently 100%, test accuracy starts off random, only showing sharp increases with increased training set size, revealing these models overfit completely at small training set sizes.

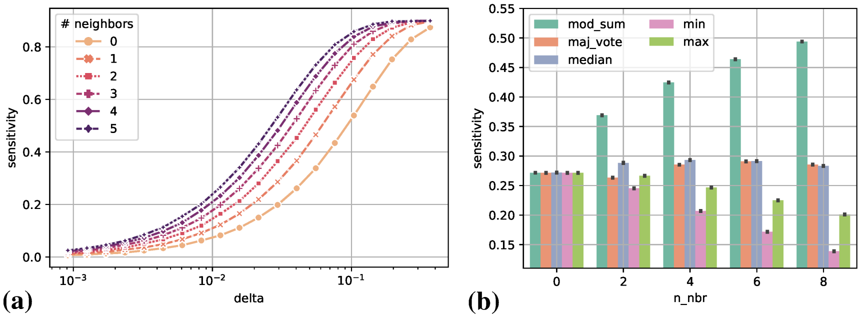

Figure 10: **Estimated noise sensitivity**. We estimate the noise sensitivity by sampling 10,000 random uniform bit sequences, and report the average estimation over 10 runs. **(a)** shows $\mathrm{NS}_\delta[f]$ for $\delta \in [e^{-7}, e^{-1}]$, for the target functions f with *mod_sum* aggregation over different neighborhood sizes. **(b)** shows the average noise sensitivity over the same value range of δ across a range of different aggregation choices.

Additional details are in the Appendix.

6 Investigating Effects of Model Architecture and Massive Dataset Sizes

In the previous section, we studied the empirical performance of `MLP` architectures on PVR tasks of different functional complexities. But the similarities between vectorized PVR and language tasks — both consisting of a sequence of tokens — suggest investigating alternate architectures, such as the Transformer [32] and even the recently proposed MLP-Mixer [30]. In this section, we examine performance variations across different architectures, as well as massive training dataset sizes.

In Figure 11 we show test performances of different architectures as training dataset size and particularly complexity are increased. We observe that `MLP` and `MLP 2x` have similar performance, while the `Transformer`

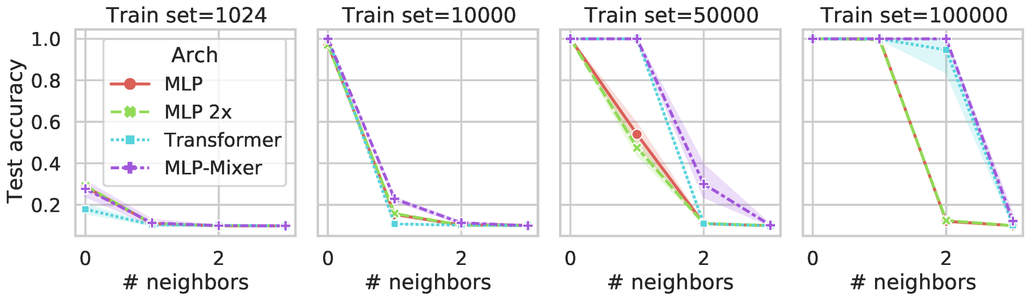

Figure 11: **Evaluating vectorized PVR tasks of different complexities demonstrates that `Transformer` and particularly `MLP-Mixer` have stronger performance.** We show test performance for these different architectures across different training set sizes and functional complexities, finding that `Transformer` and `MLP-Mixer` show significantly stronger performance, requiring fewer samples for more complex tasks, despite having more parameters. This suggests they have helpful inductive biases.

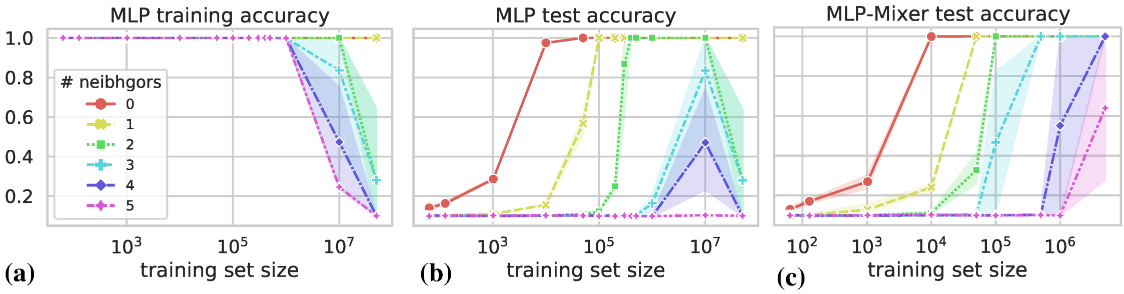

Figure 12: **Training (a) and test (b) performance `MLP`, test performance of (c) `MLP-Mixer` at massive training set sizes.** Both models show ability to tackle more complex tasks with increased training data, though `MLP` saturates on capacity, struggling to fit the training dataset at 10^7 datapoints. `MLP-Mixer` sees consistent improvements and is able to tackle tasks of complexity m at $5\text{x}10^{m+3}$ datapoints, and could achieve perfect training accuracy (not shown in the plot) for all cases tested here. Though `MLP-Mixer` runs show greater variance on more complex tasks: sometimes the optimization fails. We find that at high complexities some runs *fail to learn*, and *we ignore runs where the training accuracy is below 20%.*

and particularly `MLP-Mixer` demonstrate significantly better performance — solving tasks of higher complexity with much less data. These architectures also have larger number of parameters than `MLP` (Appendix Section A), so their better sample complexity arises from better inductive biases. This suggests a future direction of comparing representations across architectures.

Motivated by the increases in performance from dataset size in Figure 9 and to test the limits of neural network learning, we look at training with massive dataset sizes, up to 5×10^7. We show the results for `MLP` and `MLP-Mixer` in Figure 12. For `MLP`, we observe that with 10^7 training points, it is able to learn neighborhoods of size 3 (which was at random accuracy in the smaller sweep of Figure 9), but at larger number of datapoints, struggles to fit the training set. `MLP-Mixer`, with its larger capacity and better inductive bias, shows continuing performance improvements as dataset size is increased, solving more and more complex tasks, although we see variance increase with the most complex tasks, with *some seeds failing to train* (suggesting open questions on learning dynamics). Note for tasks of complexity m, we see 100% test accuracy with `MLP-Mixer` at $5 \times 10^{m+3}$ training examples.

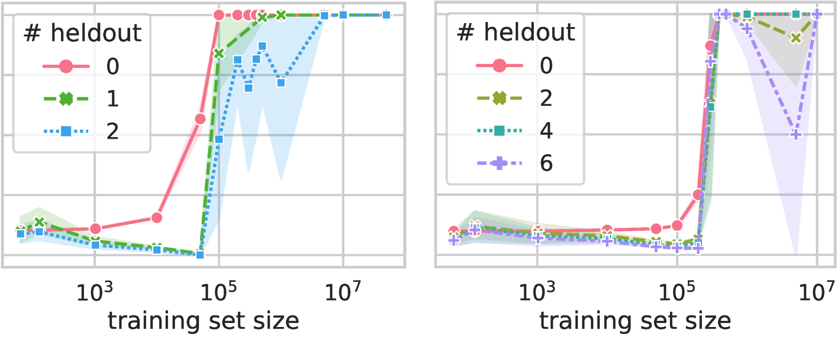

Figure 13: **A holdout shift experiment at higher complexity shows neural networks can generalize to truly unseen instances — if they don't fail at training**. We train neural networks on task complexities $m = 1, 2$ with up to $(m + 1)!$ permutations of $(0, 1, \ldots, m)$ heldout from the window of values at training, and test on *only* examples with $(0, 1, \ldots, m)$ in the values window. For $m = 3$, training is often unstable, and we discard runs with training accuracy less than 60%, with some failed runs shown in Figure 16. Strikingly, when training succeeds, the models achieve 100% accuracy on the holdout test set, strongly suggesting they are learning to reason.

7 Does high test accuracy correspond to learning reasoning?

The results of Figures 9, 12 illustrate that larger datasets help reach high test performance, with models often showing a sharp jump in test accuracy as the dataset size is increased. This raises a key question: is the high test accuracy truly indicative of learning reasoning, or for a task of complexity m, are the models simply memorizing the 10^{m+2} digits corresponding to all the different choices of (i) pointer (ii) neighborhood digits (iii) the associated label? To answer this, we devise a test by combining distribution shift and functional complexity, which provides surprising insights on the reasoning capability of neural networks.

Specifically, for a task of complexity m, we ensure the sequence $(0, 1, \ldots, m)$ never occurs in the value window in the training set. Meanwhile, the test set is adversarially constructed to contain only $(0, 1, \ldots, m)$ in the value window, but the pointers and other digits outside the value window are random. We call this task `holdout-1`. Similarly, `holdout-2`,...,`holdout-i`, etc. are constructed to hold out not only $(0, 1, \ldots, m)$, but also i permutations of it. This is the natural extension to complexity m of the Holdout Shift experiment described in Sections 3.3. But as the value is not directly the output (as in complexity 0), there is potential to avoid the systematic failure of Section 4.

The results of holding out up to $(m + 1)!$ permutations for $m = 1, 2$ and evaluating on the adversarial test set are shown in Figure 13. We observe a partial remarkable success: there are some random seeds where the neural network fails to train (not included in the plot). But in seeds where it does learn, it generalizes even when *all* permutations are held out, strong evidence that it is truly learning to reason. Some examples of failed runs are shown in Figure 16: we consistently observe a difference between *slow and fast learning*. When training converges slowly in a random seed, the test accuracy is poor, while fast training convergence results in 100% test accuracy. Together, these results suggest rich open directions on understanding task training dynamics and the mechanisms of learning.

8 Discussion

In this paper, we propose a novel benchmark, Pointer Value Retrieval, to test the limits of neural network generalization. Although PVR consists of a diverse family of tasks, with varying inputs and difficulties, they are unified by a simple rule where part of the input acts as a *pointer* indicating another part of the input which forms the *value*. This combination of positions, values and the simple rule allow us to develop nuanced tests of generalization, and our detailed empirical study demonstrates subtle systematic errors on distribution shifts, performance variations when increasing functional complexity and the effect of training set size. Our benchmark also reveals inductive biases across different model architectures that enables learning in high complexity regimes, albeit at the cost of stability. This latter point also relates to our final, striking observation, that when training can be made to succeed, neural networks are capable of learning some amount of reasoning, working on truly unseen instances. These results raise many open questions to explore on this benchmark, from variations of representations across architectures, to properties of the dataset, to the dynamics of learning the learning process, all which offer promising future directions for testing and understanding generalization.

References

[1] D. Bahdanau, S. Murty, M. Noukhovitch, T. H. Nguyen, H. de Vries, and A. Courville. Systematic generalization: What is required and can it be learned? In *ICLR*, 2019.

[2] D. Barrett, F. Hill, A. Santoro, A. Morcos, and T. Lillicrap. Measuring abstract reasoning in neural networks. In *International Conference on Machine Learning*, pages 511–520. PMLR, 2018.

[3] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[4] J. Cai, R. Shin, and D. Song. Making neural programming architectures generalize via recursion. In *ICLR*, 2017.

[5] X. Chen, C. Liang, A. W. Yu, D. Song, and D. Zhou. Compositional generalization via neural-symbolic stack machines. In *Advances in Neural Information Processing Systems*, 2020.

[6] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020.

[7] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.

[8] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. Qi, Y. Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. *arXiv preprint arXiv:2104.10133*, 2021.

[9] V. Feldman and C. Zhang. What neural networks memorize and why: Discovering the long tail via influence estimation. *arXiv preprint arXiv:2008.03703*, 2020.

[10] J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71, 1988.

[11] W. Han, Z. Zhang, B. Caine, B. Yang, C. Sprunk, O. Alsharif, J. Ngiam, V. Vasudevan, J. Shlens, and Z. Chen. Streaming object detection for 3-d point clouds. In *European Conference on Computer Vision*, pages 423–441. Springer, 2020.

[12] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 1225–1234. PMLR, 2016.

[13] J. Johnson, B. Hariharan, L. Van Der Maaten, L. Fei-Fei, C. Lawrence Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.

[14] D. Keysers, N. Schärli, N. Scales, H. Buisman, D. Furrer, S. Kashubin, N. Momchev, D. Sinopalnikov, L. Stafiniak, T. Tihon, D. Tsarkov, X. Wang, M. van Zee, and O. Bousquet. Measuring compositional generalization: A comprehensive method on realistic data. In *ICLR*, 2020.

[15] U. Khandelwal, O. Levy, D. Jurafsky, L. Zettlemoyer, and M. Lewis. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*, 2019.

[16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014.

[17] B. Lake and M. Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, pages 2873–2882. PMLR, 2018.

[18] B. Lake, R. Salakhutdinov, J. Gross, and J. Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011.

[19] W. Nie, Z. Yu, L. Mao, A. B. Patel, Y. Zhu, and A. Anandkumar. Bongard-logo: A new benchmark for human-level concept learning and reasoning. *Advances in Neural Information Processing Systems*, 33, 2020.

[20] M. I. Nye, A. Solar-Lezama, J. B. Tenenbaum, and B. M. Lake. Learning compositional rules via neural program synthesis. *arXiv preprint arXiv:2003.05562*, 2020.

[21] L. Oakden-Rayner, J. Dunnmon, G. Carneiro, and C. Ré. Hidden stratification causes clinically meaningful failures in machine learning for medical imaging. In *Proceedings of the ACM conference on health, inference, and learning*, pages 151–159, 2020.

[22] R. O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[23] J. Oh, S. Singh, H. Lee, and P. Kohli. Zero-shot task generalization with multi-task deep reinforcement learning. In *International Conference on Machine Learning*, pages 2661–2670. PMLR, 2017.

[24] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.

[25] M. Raghu, A. Irpan, J. Andreas, B. Kleinberg, Q. Le, and J. Kleinberg. Can deep reinforcement learning solve erdos-selfridge-spencer games? In *International Conference on Machine Learning*, pages 4238–4246. PMLR, 2018.

[26] L. Ruis, J. Andreas, M. Baroni, D. Bouchacourt, and B. M. Lake. A benchmark for systematic generalization in grounded language understanding. In *Advances in Neural Information Processing Systems*, 2020.

[27] J. Russin, R. Fernandez, H. Palangi, E. Rosen, N. Jojic, P. Smolensky, and J. Gao. Compositional processing emerges in neural networks solving math problems. *arXiv preprint arXiv:2105.08961*, 2021.

[28] D. Saxton, E. Grefenstette, F. Hill, and P. Kohli. Analysing mathematical reasoning abilities of neural models. *arXiv preprint arXiv:1904.01557*, 2019.

[29] J. Snell, K. Swersky, and R. S. Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017.

[30] I. Tolstikhin, N. Houlsby, A. Kolesnikov, L. Beyer, X. Zhai, T. Unterthiner, J. Yung, D. Keysers, J. Uszkoreit, M. Lucic, et al. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.

[31] E. Triantafillou, T. Zhu, V. Dumoulin, P. Lamblin, U. Evci, K. Xu, R. Goroshin, C. Gelada, K. Swersky, P.-A. Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096*, 2019.

[32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

[33] J. K. Winkler, C. Fink, F. Toberer, A. Enk, T. Deinlein, R. Hofmann-Wellenhof, L. Thomas, A. Lallas, A. Blum, W. Stolz, et al. Association between surgical skin markings in dermoscopic images and diagnostic performance of a deep learning convolutional neural network for melanoma recognition. *JAMA dermatology*, 155(10):1135–1141, 2019.

[34] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh. Large batch optimization for deep learning: Training bert in 76 minutes. 2019.

[35] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

[36] T. Zhuo and M. Kankanhalli. Solving raven's progressive matrices with neural networks. *arXiv preprint arXiv:2002.01646*, 2020.

A Dataset and Architecture Details

PVR Datasets: In the main text, we perform experiments on:

- **visual block style MNIST PVR IID dataset** For this dataset, which is used in Section 4, for the training set, we sample each position iid from the 60000 MNIST training examples. For the testset, we similarly sample each position iid from the 10000 MNIST test set. To introduce some more variation, each block position is of size 40x40, and we randomly jitter the 28x28 MNIST digit within the block.
- **visual block style MNIST PVR Holdout Shift dataset** Here for the training set, we again sample from the MNIST training set, but holding out some digits at some positions. So for the top right, we don't sample from digits $1-3$ (but sample iid from the remaining MNIST digits), for the bottom left, we don't sample from $4-6$ (but iid from all other MNIST training digits) and for the bottom right we don't sample from MNIST digits that are $7-9, 0$. The pointer is sampled iid from all MNIST training digits. For the test set we sample from the MNIST test digits, but again holding out some digits, i.e. we *only* sample iid from MNIST test digits $1-3$ for the top right corner, and so on. We also jitter each digit like above.
- **vectorized PVR dataset** For the vectorized PVR dataset, we simply sample iid from $0-9$ for each position in the sequence.

For the visual PVR tasks, specifically the MNIST visual PVR task, we use the standard `ResNet18` and `VGG11bn`, from the `torchvision` models library. We train for 10 epochs with batch size 50. We use the Adam optimizer with learning rate 10^{-4}.

For the vectorized PVR tasks, we use the following network architectures. Table 1 summarize the number of parameters in each model we use.

- `MLP`: We use an embedding layer with vocabulary size 10 and embedding dimension 64 to map each of the input token to a vector representation. The concatenated representations from all input tokens are then passed through 4 fully connected layers with output dimension 512, 1024, 512, and 64, respectively. We apply ReLU activation function after each fully connected layer. Finally, a linear layer with output dimension 10 is attached as the classifier.
- `MLP` $2\times$: The same as `MLP`, except that the output dimensions of the 4 fully connected layers are doubled: 1024, 2048, 1024 and 128.
- `Transformer`: We use the encoder part of standard Transformer [32], which consists of multiple *transformer layer*, where each transformer layer consists of a multi-head self-attention block and a MLP block. In particular, we use embedding dimension 512, 4 transformer layers, with 4 heads for each self-attention block and hidden dimension 1024 for each MLP block. Following Dosovitskiy et al. [6], we prepend a virtual *class token* with learnable vector representations, and in the final encoder output, attach a linear classifier to the representation of that token for classification.
- `MLP-Mixer`: MLP-Mixer is a new architecture recently proposed by Tolstikhin et al. [30]. It is similar to Transformers, except that the multi-headed self-attention layers are replaced with fully connected layers. We use embedding dimension 512 and 4 *mixer layers*, with token-MLP dimension 768 and channel-MLP dimension 2048. Similar to `Transformer`, we use a *class token* for the purpose of classification.

We did a small scale hyperparameter sweep on different optimizers (SGD, Adam [16], Adagrad [7], LAMB [34]), learning rates ($1 \times 10^{-3} \sim 4 \times 10^{-1}$) and batch sizes (128, 256, 512, 1024, 2048, 4096). In the end we choose the following hyperparameters by balancing performance and stability across different setups: We use weight decay 1×10^{-5}, SGD optimizer with momentum 0.9, and cosine learning rate scheduling with base learning rate 0.05, and a linear warmup period of 10 epochs. The batch size is 1024, and we train for 200 epochs. For studies with tiny training sets (e.g. 64), we train for at least 800 iterations. The training was done with NVidia P100 / V100 GPUs.

Table 1: Number of parameters for models used in the empirical studies.

Architecture	Parameter count
`MLP`	1,445,194
`MLP` 2×	5,052,426
`MLP-Mixer`	8,495,674
`Transformer`	8,429,066

B Additional Results from Initial Experiments on Visual PVR and Distribution Shift

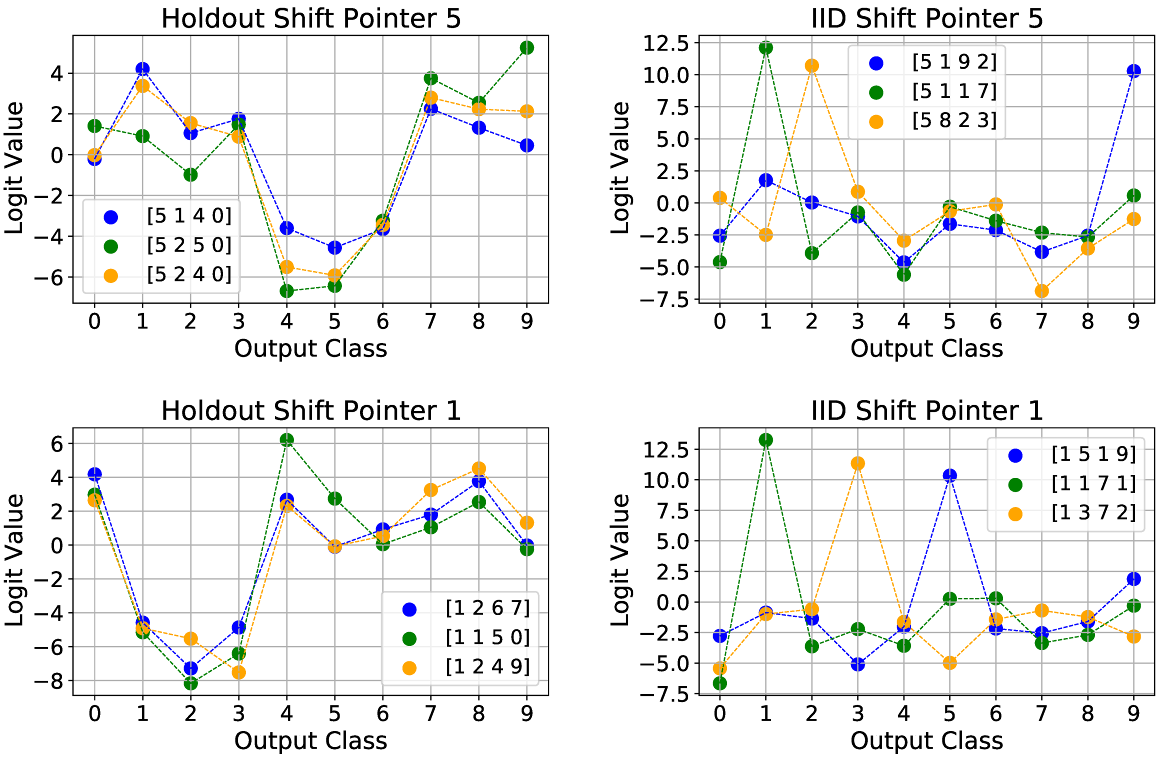

Figure 14: **Analyzing the failure of neural networks on Holdout shift reveals systematic mistakes from learned correlations between pointer digits and labels.** Additional logit results, compare to Figure 4 in the main text.

C Noise Sensitivity Analysis

We applied *noise sensitivity* of boolean functions in the main text to characterize the complexity of our vectorized PVR tasks with different neighborhood sizes and aggregation functions. Specifically, for a function f with boolean (binary bits) inputs, and $\delta \in [0, 1]$, the *noise sensitivity* $\mathrm{NS}_\delta[f]$ measures the probability that $f(x) \neq f(y)$, where x consists of uniformly random bits, and y is formed from x by flipping each bit independently with probability δ. It is widely used to measure the stability and closely related to various complexity measure of boolean functions. For example, for $\delta \in (0, 1/2]$, the Fourier spectrum of a binary output boolean function f can be shown to be ϵ-concentrated on degree up to $1/\delta$ with $\epsilon \leq 3\mathrm{NS}_\delta[f]$ [22,

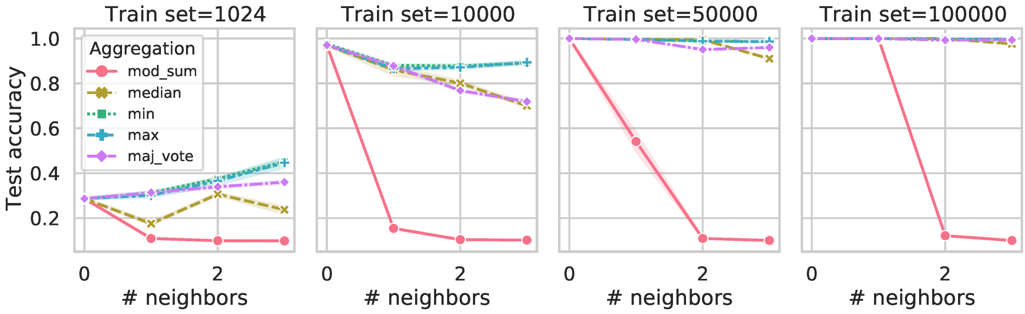

Figure 15: **Evaluating different aggregation functions for PVR tasks shows *mod_sum* is the most challenging.** We show test performance for different aggregation functions across varying dataset size and functional complexity. The empirical results support the intuitive observation that *mod_sum* is the most challenging.

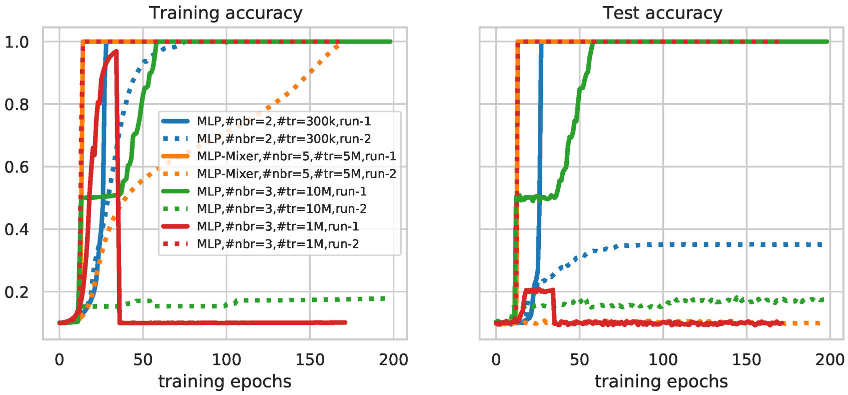

Figure 16: Learning curve examples. Each color show one pair of experiments with identical hyperparameters except different random seeds, demonstrating drastically different behaviors under identical setting.

Proposition 3.3].

We extend it to measure the complexity of our target functions by representing the inputs of our vectorized PVR task as bit vectors. In particular, each digit $x \in \{0, \ldots, K-1\}$ in our input sequence is represented as D bits by the standard binary representation of unsigned integers, where D is the smallest integer such that $K \leq 2^D$. When $K \neq 2^D$, a subset of random uniform bit sequences will fall out of the valid range of x. We simply extend the definition of the target functions to take digits with arbitrary values, but convert them into the "valid range" via mod K as a preprocessing.

Figure 15 illustrates test accuracy for different aggregation functions across varying PVR functional complexity and dataset sizes. The learning results are consistent with our measurements of task complexity using noise sensitivity: the tasks become noticeably more challenging to learn as the neighborhood size increases only when using the *mod_sum* aggregation function.

D Fast Learning vs Slow Learning

We present a few cases of training curves in Figure 16. Each color show a pair of runs with identical hyperparameters and task setups, but with different random seeds. We found that learning could become unstable when the task complexity increases. For example, the red pair of experiments show that one of the training job suddenly collapse during training and fail to to recover from that in the remaining epochs, while another run with a different random seed succeeds and generalize perfectly. Similarly, the green pair show one run that never reaches 20% *training* accuracy, but another run with perfect *training and test* accuracy.

A more interesting pattern we observe is "slow learning" vs "fast learning": sometimes, as demonstrated by the blue pair and the orange pair, both runs succeed in fitting the training set perfectly, yet they have completely different generalization capability. We found that when a model converges rapidly on the training set, it usually generalize well; on the other hand, when a model converges slowly and smoothly on the training set, it usually generalize poorly. It seems in the latter case, the network is just slowly memorizing the training examples without learning the actual concept. This is consistent with earlier theoretical work that quantify generalization via training speed [12]. However, it remains open question to identify here what is the factor that cause the two networks under identical setting to behave differently.