
On the creation of narrow AI: hierarchy and nonlocality of neural network skills

Eric J. Michaud^{1,3*} Asher Parker-Sartori² Max Tegmark^{1,3}

¹ Department of Physics, Massachusetts Institute of Technology

² Department of EECS, Massachusetts Institute of Technology

³ The NSF AI Institute for Artificial Intelligence and Fundamental Interactions

Abstract

We study the problem of creating strong, yet narrow, AI systems. While recent AI progress has been driven by the training of large general-purpose foundation models, the creation of smaller models specialized for narrow domains could be valuable for both efficiency and safety. In this work, we explore two challenges involved in creating narrow AI systems, having to do with basic properties of how neural networks learn and structure their representations. The first challenge regards when it is possible to train narrow models from scratch. Through experiments on a synthetic task, we find that it is sometimes necessary to train networks on a wide distribution of data to learn certain narrow skills within that distribution. This effect arises when skills depend on each other hierarchically, and training on a broad distribution introduces a curriculum which substantially accelerates learning. The second challenge regards how to transfer particular skills from large general models into small specialized models. We find that model skills are often not perfectly localized to a particular set of prunable components. However, we find that methods based on pruning can still outperform distillation. We investigate the use of a regularization objective to align desired skills with prunable components while unlearning unnecessary skills.

1 Introduction

Today, the most competent AI systems in *any* particular domain are general systems that are relatively competent in *every* domain. The best models at math and coding are also broadly knowledgeable about a very diverse array of topics, from Roman history to home cooking recipes to medical diagnostics. And when domain-specific models are created today, they are typically general foundation models fine-tuned on a particular task, rather than new models trained from scratch [1, 2], though with some notable exceptions [3]. This state of affairs is convenient and powerful, since a single general model can be used for a variety of applications [4].

However, generality has downsides for both efficiency and safety. For instance, AI systems used as coding assistants possess a large amount of knowledge which is never needed in those applications. Instead, we might like to use smaller, specialized networks which preserve the coding knowledge of general systems without the same breadth of irrelevant knowledge. Narrow systems may also pose fewer safety risks than general systems. For instance, narrow systems may have fewer dangerous capabilities that pose CBRN risks [5], be easier to understand mechanistically [6, 7], or be easier to verify properties of [8–10]. More imaginatively, for systems to operate autonomously in the world requires a large breadth of skills, and an ecosystem of narrow “tool AI” systems may therefore reduce loss-of-control risks and better support human agency over the long term [11].

In this work, we investigate some basic questions about how neural networks *learn* and *implement* skills that are relevant to the problem of creating narrow AI systems. As summarized in Figure 1,

* ericjm@mit.edu

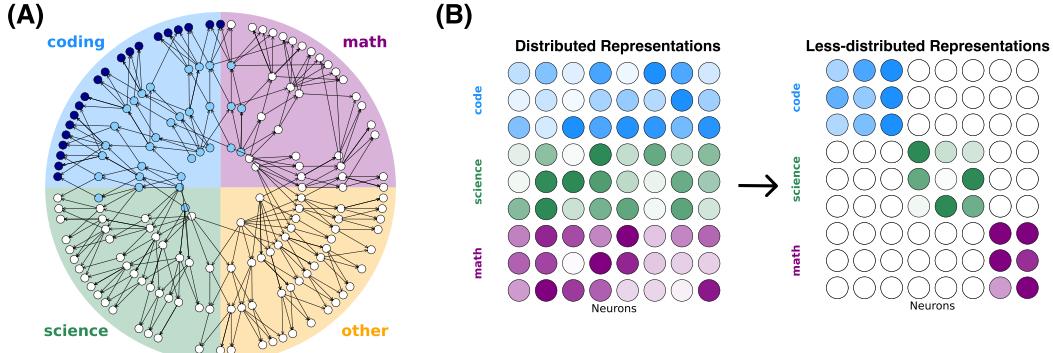


Figure 1: We study two challenges to making strong, narrow-purpose AI models. **(A)**: Data may have *hierarchical* structure. If skills have a hierarchical dependence, where some skills are only learnable after more primitive skills are learned first, then it sometimes may be necessary to train on a broad distribution of data to learn certain narrow skills within that distribution. These dynamics may mean that general-purpose models must be trained to achieve good performance on some domains. **(B)**: Model features are *distributed*. By default, skills may not be localizable to a particular set of model components (e.g. neurons). In this case, pruning of model components won’t precisely retain wanted skills and remove unwanted skills from models. We explore methods for aligning the model features relevant to particular domains with a smaller subset of model components while unlearning others.

we focus on two main themes. First, we consider the question of when it is possible to train well-performing neural networks from scratch on a narrow data distribution. Through experiments on a novel synthetic task with *hierarchical structure*, we find that it can be necessary to train networks on a broad distribution to efficiently learn narrow tasks within that distribution. These results contribute to a growing literature on how task structure influences neural network learning dynamics [12–15], but with special relevance to the problem of creating narrow AI. Second, we consider the question of whether one can use pruning to turn broad networks into smaller narrow ones. We find that the *nonlocality* of network representations to prunable model components poses a challenge for this goal. While distributed representations have been extensively studied in the context of neural network interpretability [16] and in the classical connectionist literature [17, 18], we study how this property of neural network computation impacts pruning and unlearning [5] for creating narrow AI. Our specific contributions are as follows:

- We describe a synthetic task, *compositional multitask sparse parity* (CMSP), extending the multitask sparse parity task of Michaud et al. [19]. We find that networks trained on this task exhibit extremely strong curriculum learning effects, where it is necessary to train on a broad distribution of tasks in order to learn certain other tasks.
- We study pruning and unlearning on our networks trained on CMSP. We observe that tasks are often distributed and distinct subtasks are entangled, making pruning an imperfect strategy for “narrowing” the breadth of network skills. However, we find that a simple group-sparsity regularization objective can be used to sparsify networks and unlearn skills.
- We perform an empirical comparison of methods for creating narrow systems on MNIST and in LLMs. We tentatively find that methods based on pruning outperform distillation and training networks from scratch for the creation of smaller, more narrow systems.²

This work is organized as follows. In Section 2, we briefly describe methods. In Section 3 we perform a detailed case study of training and pruning on the CMSP task. We define the task (Section 3.1), study network training dynamics on it (Section 3.2), find distributed representations in these networks (Section 3.3), and use a regularization method for pruning and unlearning in these networks (Section 3.4). We then compare various methods for creating narrow models, studying MNIST in Section 4 and language models in Section 5. We discuss related work in Section 6 and conclude in Section 7.

²Code for the experiments in this paper can be found at: <https://github.com/ejmichaud/narrow>.

2 Methods

Pruning: We aim to preserve the performance of a model $f(\cdot; \theta)$ on a distribution \mathcal{D}_N while pruning model components. Let g be a collection of parameter indices i corresponding to prunable components of the model (e.g. the in-weights and out-weights of a neuron), and let G denote the collection of all such groups. After ablating a specific g , we denote the new parameters θ_g^* . To perform pruning, for each $g \in G$, we compute an *ablation score* $s_g = |\mathbb{E}_{(x,y) \in \mathcal{D}_N} [L(f(x; \theta), y) - L(f(x; \theta_g^*), y)]|$, the absolute change in the model’s expected loss L after pruning g . We sort groups by their ablation score and prune greedily to the desired sparsity from lowest to highest ablation score. Where feasible, we estimate s_g empirically by manually ablating the group across many samples from \mathcal{D}_N . When this is computationally intractable, we instead use a linear estimate which we refer to as an *attribution score* after [20, 21]: $\hat{s}_g = |\sum_{i \in g} \frac{\partial L}{\partial \theta_i}(-\theta_i)| = |\frac{\partial L}{\partial \theta} \cdot (\theta_g^* - \theta)|$, where $\partial L / \partial \theta$ is the model’s gradient on the distribution \mathcal{D}_N , which can be computed once and reused for all $g \in G$.

Regularization: We also experiment with making networks more prunable by performing additional training with a “group lasso” regularization penalty on the model weights [22–25]. The group lasso penalty R is the L1 norm of the L2 norms of each parameter group: $R(\theta) = \sum_{g \in G} \sqrt{\sum_{i \in g} \theta_i^2}$. This penalty incentivizes the weights to become sparse at the level of entire groups g . When we perform “group lasso training” on a distribution \mathcal{D}_N , we minimize the loss $\mathbb{E}_{(x,y) \sim \mathcal{D}_N} [L(f(x; \theta), y) + \lambda R(\theta)]$.

Distillation: We also explore distilling knowledge from a teacher model using the standard algorithm employed in [26], minimizing the KL divergence between student and teacher output distributions.

3 Case study: compositional multitask sparse parity

In this section, we conduct a detailed study of both curriculum learning and pruning on simple synthetic task, which we call *compositional multitask sparse parity* (CMSP).

3.1 Defining compositional multitask sparse parity (CMSP)

The compositional multitask sparse parity (CMSP) task is a simple extension of the *sparse parity* task recently studied in [27] and the *multitask sparse parity* task studied in [19], described below.

Barak et al. [27] studied neural network training dynamics on the *sparse parity* (SP) task. This is a binary classification problem on binary strings x of length n , where the label of a given sample x is the parity of the bits at a subset I of k indices: $y = \oplus_{l=1}^k x[I_l]$. Strings x are sampled uniformly. Barak et al. observed that the loss curve for neural networks trained on this task exhibits a sharp drop after an initial plateau, a case of the sort of “emergence” which has been observed in LLMs [28, 29] and in other toy settings [30].

Michaud et al. [19] extended the task to *multitask sparse parity* (MSP). In the MSP problem, input strings consist of $m + n$ bits, where the first m bits are called “control bits” and the last n bits are called “task bits”. These leading m control bits encode which “subtask” must be solved in each problem instance. Instead of a single length- k set of indices I , a collection of m sets of task bit indices (typically of equal length k) is chosen $\{I_1, \dots, I_m\}$. For each sample, only one control bit is ON (1) while the others are OFF (0). If control bit t is ON, then that sample’s label is the parity of bits I_t : $y = \oplus_{l=1}^k x[(I_t)_l]$. Michaud et al. [19] found that when subtasks are power law distributed in frequency, where the probability that control bit i is ON is $p_i \propto i^{-(\alpha+1)}$, the mean loss exhibits power-law scaling [31, 32] while individual subtasks are learned at different times proportional to their frequency [33]. They conjecture that LLM learning dynamics may be similar.

One of the main limitations of the MSP task, and the associated model of neural scaling from [19], is that subtasks are independent. However, the world, and the problem of learning from it, intuitively has a hierarchical structure—in order to learn to do certain concepts and tasks, we must first learn simpler concepts and tasks. Accordingly, humans are taught in a curriculum, where simpler concepts and tasks are learned first, and then composed into more complex ones later. To capture this structure, we now introduce the toy task *compositional multitask sparse parity* (CMSP). We show some CMSP samples to the right:

CMSP Samples: m=3, n=9, k=3		
	x	y
Atomic	100101001010	0
Atomic	100010100011	1
Atomic	010110111110	1
Atomic	001101101110	0
Compositional	110001101011	1
Compositional	110111001010	0
Compositional	011110001110	1
Compositional	111000110100	1
Control Bits		
Task Bits		

CMSP is similar to MSP, except that (1) we require subtask indices to be disjoint: $I_i \cap I_j = \emptyset$ if $i \neq j$, and (2) multiple control bits can now be ON at the same time, in which case the label is the parity of the bits in the union of the indices for each subtask. If control bits i, j are both ON, the label is the parity of the task bits at indices $I_i \cup I_j$. We call samples for which only one control bit is ON “atomic” and samples for which multiple control bits are ON “composite”. So if $k = 3$ for all subsets I_i , then on atomic samples networks must compute the parity of 3 input bits, on composite samples with two control bits ON the label is the parity of 6 input bits, and so on. Above, we illustrated some CMSP data samples with $m = 3$, $n = 9$, and $k = 3$.

3.2 Learning dynamics on CMSP

We find that neural network training on CMSP exhibits extremely strong curriculum learning effects. To denote different types of CMSP samples, we list the ON control bits for those samples. So, given a choice of m, n, k and I_1, \dots, I_m , $\mathcal{D}_{\{0\}}$ denotes all samples for atomic subtask 0, and $\mathcal{D}_{\{0,1,2,3\}}$ denotes all samples from the composite subtask where the first four control bits are ON. For each subtask, there are 2^n possible samples in that subtask, since the m control bits are fixed but the n task bits are free.

We first train ReLU MLPs with 1-2 hidden layers of width 128 with the Adam optimizer on CMSP samples with $m = 4$, $n = 64$, $k = 4$, and 2000 samples per task (atomic or composite) per batch. We show the learning dynamics of these networks in Figure 2. We find that when we train on a dataset containing both atomic and composite samples $\mathcal{D}_{\{0\}} \cup \mathcal{D}_{\{1\}} \cup \mathcal{D}_{\{2\}} \cup \mathcal{D}_{\{3\}} \cup \mathcal{D}_{\{0,1,2,3\}}$ in equal proportion, atomic tasks are learned before composite ones.

Something interesting happens when we remove atomic samples from the dataset: we find that learning composite tasks takes dramatically longer. When we train on $\mathcal{D}_{\{0\}} \cup \mathcal{D}_{\{1\}} \cup \mathcal{D}_{\{2\}} \cup \mathcal{D}_{\{3\}} \cup \mathcal{D}_{\{0,1,2,3\}}$ with 10000 samples per batch (2000 per subtask), across 40 seeds, we find that 27/40 networks converge on composite subtask $\mathcal{D}_{\{0,1,2,3\}}$ within 2×10^9 samples, and the networks that

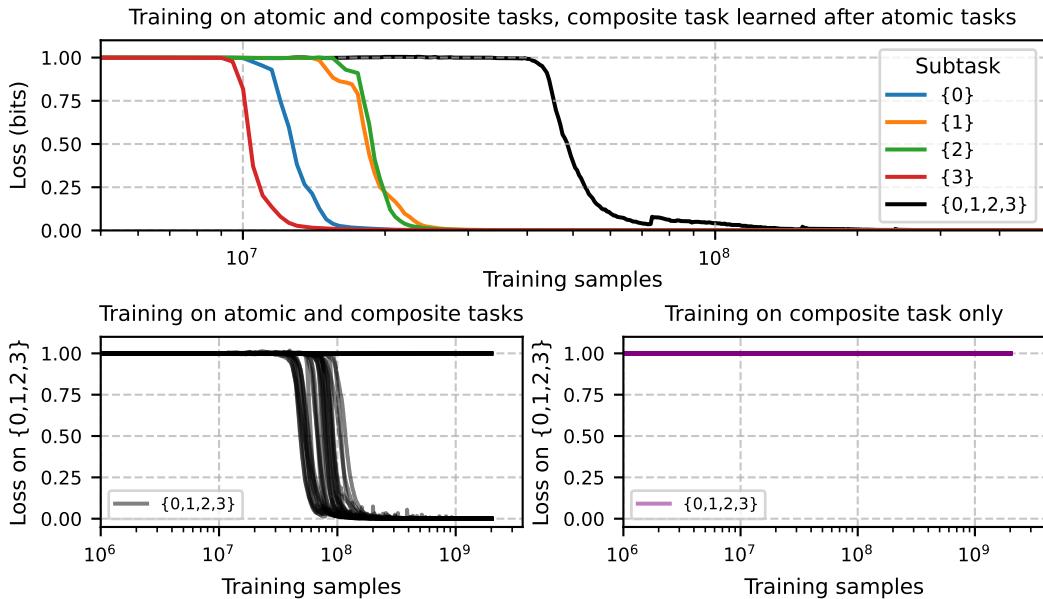


Figure 2: Training dynamics on compositional multitask sparse parity. **Top:** training dynamics for a single network trained on four atomic subtasks $\{0\}$, $\{1\}$, $\{2\}$, $\{3\}$, and their composition $\{0,1,2,3\}$. **Bottom left:** loss on compositional subtask $\{0,1,2,3\}$, when training on atomic subtasks and their composition, as in the top subplot, across 40 network seeds (depth 3, width 128). **Bottom right:** loss on compositional subtask $\{0,1,2,3\}$, when only training on samples from that composite task, *without* also training on atomic tasks, across 40 network seeds (depth 3, width 128). We find that removing the atomic tasks prevents our networks from learning the composite task. We report the minimum loss within the previous 100 steps of training to filter out loss spikes.

do converge typically converge within 2×10^8 samples (bottom left). However, when we train just on $\mathcal{D}_{\{0,1,2,3\}}$, with a batch size of 2000 samples, we find that 0/40 networks converge within 2×10^9 samples in Figure 2 (bottom right). It is much more efficient to train on a broader distribution in order to learn subtask $\mathcal{D}_{\{0,1,2,3\}}$ than just training on $\mathcal{D}_{\{0,1,2,3\}}$ on its own.

This result makes sense given the exponential hardness of learning parities [34]. Since the atomic task bits are disjoint in CMSP, networks can compute the parity of the composite tasks by computing the parity of features learned for the atomic tasks, reducing the complexity of the problem. For instance, one could imagine that depth-3 networks would compute the parity of the atomic tasks in the first layer, then the composite tasks in the next layer. Curiously though, we find that single-hidden-layer (depth-2) networks can still take advantage of a curriculum in Appendix A. We leave a mechanistic analysis of how networks learn these tasks to future work.

These sorts of dynamics may be a part of the explanation for why large-scale general-purpose models perform so strongly at many narrow, valuable tasks. We emphasize, however, that the toy task where we observe these curriculum effects is fairly contrived. It is not clear to what extent there are similarly strong effects on real-world tasks, and in many domains it is in fact possible to train narrow specialized models by only training on data in that domain. For instance, self-driving cars do not need to be trained on a broad corpus of text like LLMs.

Our work shows that for *some* types of tasks, it may be necessary to train on a broad data distribution in order to learn some subtasks efficiently. Accordingly, we now turn to the question of how to efficiently transfer knowledge from large models into smaller specialized ones.

3.3 Nonlocal representations in CMSP networks

If the circuitry for some subtasks were localized to a particular set of neurons, and the circuitry for other subtasks were localized to different neurons, then the task of specializing broad networks into narrow ones would be trivial. One could simply prune away neurons (or other model components, e.g. attention heads in transformers) associated with some subtasks, and keep others. However, often the situation seems more complicated than this ideal, which will be a focus for the rest of this paper.

A related problem has recently been studied in neural network interpretability, where it has been observed that individual computational units in neural networks, such as neurons, are *polysemantic*, activating across a wide variety of unrelated inputs [35, 36]. Accordingly, many assume that the true model “features” do not align with architectural components like neurons. Multiple explanations have been proposed for this phenomenon. One is simply that the model architecture does not always “privilege” a particular basis [37, 38], though other incidental reasons for polysemanticity have also been proposed [39]. Another explanation of polysemanticity is the *superposition hypothesis* [40, 36]. The superposition hypothesis suggests that the need to represent more features than there are dimensions or neurons prevents features from being represented as orthogonal directions in the feature space, and therefore all features cannot be aligned with standalone model dimensions or neurons. Recently, studies that use sparse autoencoders to identify monosemantic model features have found that most features are highly distributed across a large number of dimensions of activation space [16].

In our CMSP networks, we observe a related problem. By default, without any explicit regularization, there is no incentive for the network to localize certain circuits cleanly into a particular set of neurons. We train networks with two hidden layers on a dataset of CMSP samples with $m = 6, n = 18, k = 3$ with two different skill trees: $\{0\}, \{1\}, \{2\}, \{0,1,2\}$, and $\{3\}, \{4\}, \{5\}, \{3,4,5\}$ until convergence. In Figure 3, we visualize the connectivity of a 2-hidden-layer MLP trained on this dataset, and see that the network is densely connected without obvious structure.

We attempt to prune this network to retain performance on subtask $\{0,1,2\}$ while unlearning $\{3,4,5\}$. With MLPs, our groups of parameters g are the in-weights, bias, and out-weights for each hidden layer neuron, so that $|G|$ is the total number of hidden neurons. We compute ablation scores on the distribution $\mathcal{D}_N = \mathcal{D}_{\{0,1,2\}}$ (estimated on 2000 samples) and prune greedily, as described in Section 2. In Figure 3, we show accuracy on subtask $\{0,1,2\}$ vs. sparsity with this pruning strategy. When applied naively, network performance tends to degrade quickly, since without explicit regularization the network is not optimized to be naively prunable (see in Appendix Figure 8 and Figure 9 for more on how ablation scores for each subtask vary across neurons). However, when we perform an additional 1000 steps of training on $\mathcal{D}_{\{0,1,2\}}$ (5000 samples per batch) after pruning (after either removing neurons from the architecture or pinning pruned weights at zero) we can recover

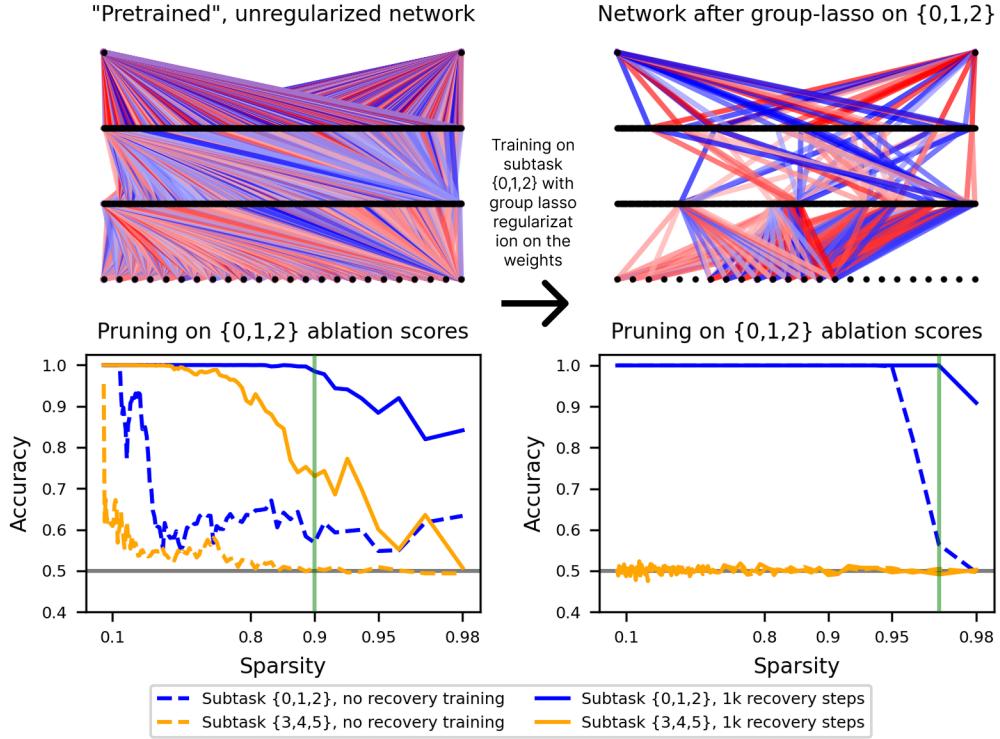


Figure 3: **Top:** We visualize the connectivity of 2-hidden-layer MLPs trained on a CMSP distribution with subtasks $\{0\}$, $\{1\}$, $\{2\}$, $\{0,1,2\}$, $\{3\}$, $\{4\}$, $\{5\}$, $\{3,4,5\}$, visualized before (**left**) and after (**right**) regularizing network weights with the group lasso sparsity penalty while training on subtask $\{0,1,2\}$. We find that network connectivity becomes sparse after regularizing. Negative weights are shown in blue and positive ones in red, with width proportional to the norm of the weight. **Bottom:** we show how pruning affects task performance on subtasks $\{0,1,2\}$ and $\{3,4,5\}$ at varying sparsity levels. We prune neurons based on the absolute change that ablating them has on the loss on subtask $\{0,1,2\}$. We find that subtasks here are nonlocal and entangled in the “pretrained” network (**left**). As we prune neurons according to their relevance on subtask $\{0,1,2\}$, at the sparsity at which performance on subtask $\{0,1,2\}$ accuracy drops below 98% (green line), we can still recover some performance on subtask $\{3,4,5\}$ with a small amount of additional training. Thus naive pruning here has not completely and robustly unlearned subtask $\{3,4,5\}$. However, after regularizing the weights (**right**), we find that not only we can more aggressively prune the network, but we have also robustly unlearned subtask $\{3,4,5\}$. Note that the degree to which subtasks are nonlocal and entangled in the pretrained networks depends on seed and width, and we show a variety of additional curves in Figure 10.

performance at higher sparsity levels. Since we are training just on composite samples, 1000 steps is not enough to re-learn this task on its own, so if we can recover performance, that will be because the mechanisms for the task were somewhat preserved after pruning. We observe that often, the tasks $\{0,1,2\}$ and $\{3,4,5\}$ are *entangled* – when we prune as aggressively as we can while being able to recover performance on subtask $\{0,1,2\}$, we can often still recover some performance on subtask $\{3,4,5\}$. However, for our CMSP networks the degree of entanglement is highly seed-dependent, and sometimes the subtasks are disentangled enough that pruning as aggressively as possible on one subtask does in fact robustly unlearn the other. We show pruning curves across seeds and widths in Appendix Figure 10.

3.4 Regularizing to “narrow” networks

We find that we can use a simple regularization penalty to simultaneously unlearn some tasks while incentivizing the network to move its features to be less distributed, allowing for more aggressive pruning. As described in Section 2, we simply perform additional training on $\mathcal{D}_{\{0,1,2\}}$ with a group

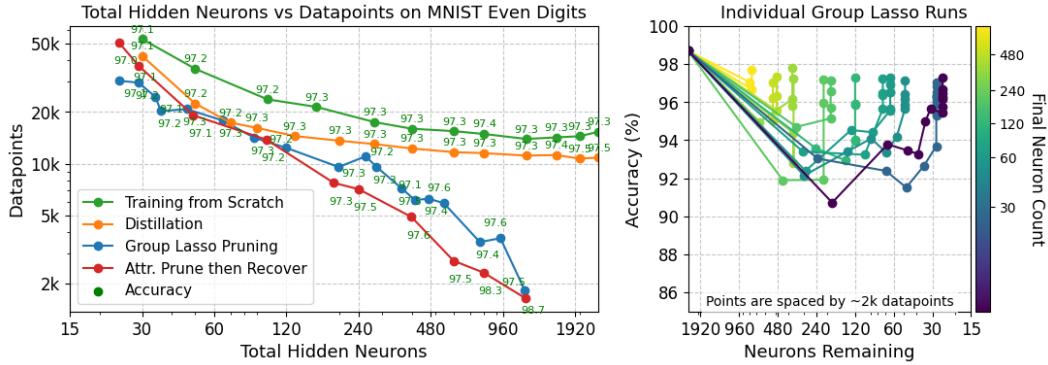


Figure 4: **Left:** We compare the performance of distillation, training from scratch, and two pruning approaches for creating small networks that classify MNIST even digits. Pruning-based approaches Pareto-dominate distillation, achieving high compression ratios with fewer datapoints. All points are averaged over 10 individual training runs. **Right:** When pruning using group lasso, it often helps to first prune rapidly, degrading performance, and then recover performance with no regularization. Each line represents a single training run, with a new point logged every 2,000 datapoints. Darker lines correspond to more aggressive pruning runs that attain lower final neuron counts.

lasso sparsity penalty on the network weights. We aim for this penalty to “clean up” circuitry [41] not relevant to prediction on $\mathcal{D}_{\{0,1,2\}}$ while also sparsifying the weights across groups of parameters, allowing for easier pruning.

We apply this regularization to our CMSP networks, training with penalty strength $\lambda = 10^{-3}$ with a batch size of 2000 for 10000 steps. With this regularized network, we then re-compute neuron ablation scores and prune. In Figure 3 we find that, in CMSP networks, this method is effective at unlearning skills and making the network more prunable. We find that we can prune more aggressively while retaining performance on subtask $\{0,1,2\}$, and we are not able to recover performance on $\{3,4,5\}$ at any sparsity level. These results complete our case study on CMSP: in order to efficiently learn compositional tasks, we train on a broad distribution of tasks. Then, we can regularize the network to turn this broad network into a much smaller one which nevertheless retains the broad network’s performance on the target compositional task.

4 Pruning vs distillation: MNIST

We now consider the problem of creating narrow systems in more natural domains, first on MNIST. As our narrow subtask, we choose only even digits from the original MNIST dataset [42]. We compare the resources required to achieve good performance on this narrow task when (1) training from scratch on the narrow task, and (2) distilling models from a general teacher on this task, (3) using group-lasso regularization to prune a large general model and (4) using attribution-based pruning and then recovery training on a large general model. Our teacher model is a ReLU MLP with two hidden layers each of width 1200, as in [26], and achieves 98.7% accuracy on the test set. When pruning, we use this same teacher model as our initial model and prune its hidden neurons to create a smaller network.

When using distillation (2), we use the approach of Hinton et al. [26] with $T = 20$. When pruning with the group lasso penalty (3), we use λ values ranging from 0.001 to 0.008. Unlike in Section 3.4, we regularize and prune simultaneously, pruning neurons when their L2 norm drops below 0.05. When the number of remaining neurons drops below a target threshold, we remove the pruning penalty and continue training to recover lost performance during pruning. When we prune up front and then separately recover performance (4), we use attribution scores as described in Section 2.

To compare methods, we require that each method reach a test-set accuracy of 97 percent, and we then plot the frontier of neuron count versus datapoints subject to that threshold. As seen in Figure 4 (right), it is often optimal to first aggressively prune the network down to the desired size and further train it until it reaches the requisite accuracy.

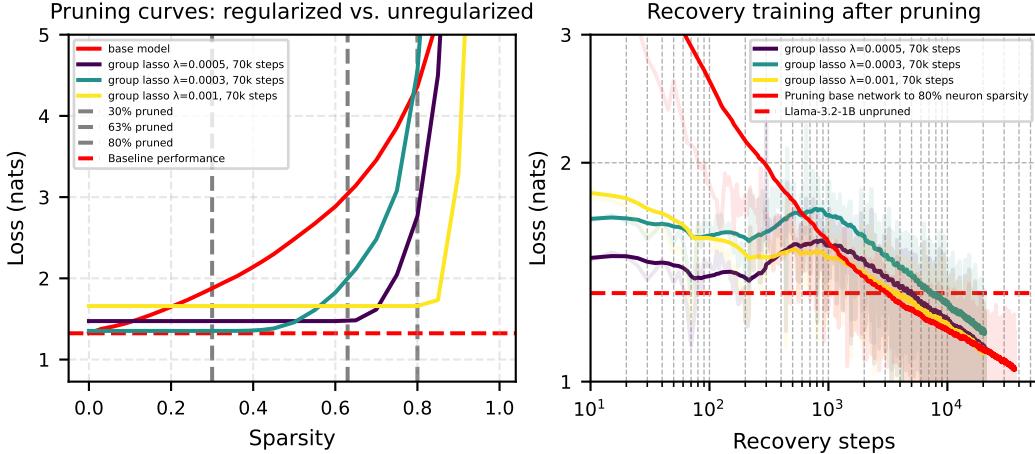


Figure 5: **Left:** Neuron sparsity vs. loss curve for LLMs finetuned with group lasso regularization with varying λ for 70k steps, as compared to the base network. Regularization flattens the sparsity vs. loss curve, at the cost of slightly degrading model performance. **Right:** After pruning our networks to 30%, 63%, and 80% sparsity for our runs with λ of 5e-4, 3e-4, and 1e-3, respectively, we recover performance with additional training. We find that we can recover performance lost during pruning, including in the network that was pruned without first using group-lasso regularization training. The plot shows an exponential moving average of batch losses, with individual batch losses shaded in back.

We find that while group lasso pruning is highly sensitive to choices in hyperparameters, both pruning methods Pareto-dominate distillation and training from scratch, especially at high neuron counts. Moreover, while other methods cannot consistently bridge the 97 percent accuracy threshold with fewer than 25 hidden neurons, aggressive pruning can consistently shrink the network’s size to a lower absolute limit.

5 Pruning vs distillation: LLMs on Python documents

We next study LLMs. As our narrow task \mathcal{D}_N , we choose next-token prediction on Python documents in the [GitHub Code Dataset](#).

We first prune the neurons in the MLP blocks of Llama-3.2-1B [43]. Later, we will also consider pruning residual stream dimensions, which involves pruning all model parameters that “read to” or “write from” a dimension of the residual stream [37]. We use attribution scores when pruning, and show that the attribution scores correlate moderately with true ablation scores in Appendix Figure 11. In Figure 5 (left), we show neuron sparsity vs loss curves. When pruning neurons naively, we see that loss increases quickly at low sparsity levels. We also experiment with applying group lasso regularization while further training on Python documents (learning rate of 2e-6, max length 512, 18 documents per batch) and find that this training does indeed level out the sparsity vs. loss curve, albeit at a slight cost to the loss. Fortunately, we find that we can recover lost performance after pruning by doing a small amount of additional training on Python documents in Figure 5 (right). We find that despite the loss increasing substantially after naively pruning, we can also quickly recover that lost performance, and overall this strategy seems to be more efficient than using group lasso training. We therefore next compare naive pruning + recovery training against distillation and training networks from scratch.

We train networks with the Llama 3 [43] architecture of varying shape and size. We use a learning rate of 5e-4, sequence length of 1024, and batch size of 64. For distillation, we use Llama-3.1-8B as a teacher with $T = 2$. For pruning + recovery, we prune Llama-3.2-1B to varying levels of neuron and residual stream sparsity, shown in Appendix Table 3. In Appendix Figure 12 we show learning curves for these three approaches. In Figure 6, we tentatively find that pruning substantially outperforms training from scratch and distilling a model from scratch on the data-parameter frontier. Given a

target narrow network size and a fixed data budget, if one already has access to a general model, it appears to be more efficient to prune that model than it is to perform distillation.

We also experiment with pruning random model components instead of components selected by attribution score and find that random pruning performs similarly [44]! In Appendix Figure 13, we find that after a moderate number of recovery steps, attribution pruning and random pruning result in the same recovered performance. This result is in line with our earlier discussion of nonlocality, and the empirical findings of Bricken et al. [16], that monosemantic features are distributed widely across model components. While some studies have found some geometric similarity between functionally similar features [45, 46], in our case it does not seem like the relevant features for our task are localized into a set of “Python” vs. “non-Python” neurons, at least that attribution pruning identifies.

Unlearning: We lastly evaluate whether neuron pruning can robustly unlearn downstream tasks. We consider three pruning methods applied to Llama-3.2-1B at sparsity levels of 30%, 63%, and 80%: (i) random neuron pruning, (ii) attribution-based pruning on Python code, and (iii) models trained with group-lasso regularization on Python documents from Figure 5. We evaluate each pruned model on four benchmarks: counterfact [47], AI2-ARC (ARC-Easy) [48], WMDP-Bio, and WMDP-Cyber [49] for hazardous knowledge. We found that the base model performed poorly on WMDP-Chem. For each model, we first measure baseline accuracy after pruning, then fine-tune on each benchmark’s training set with batch size 4, sweeping learning rates between 1e-6 and 1e-4 and reporting the best-performing configuration on the test set. In Table 1, we find that all pruning techniques, including random pruning, robustly unlearn non-Python skills with roughly similar effectiveness, except that at low sparsity, attribution pruning does not achieve much unlearning. This result contrasts somewhat with our findings on CMSP networks in Sections 3.3 and 3.4, where training with the group lasso penalty led to stronger unlearning than naive pruning.

Table 1: **Unlearning:** We prune Llama-3.2-1B using various techniques to create narrow Python-specific models, then evaluate on other tasks to assess whether pruning led to unlearning. Accuracies are reported for the pruned model before finetuning and after finetuning on a train set of each benchmark (non-FT \rightarrow FT). \checkmark indicates successful unlearning (FT score $\geq 10\%$ lower than base).

Method	Spar.	counterfact	AI2-ARC	WMDP Bio	WMDP Cyber
Base	0%	0.18 \rightarrow 0.49	0.65 \rightarrow 0.67	0.52 \rightarrow 0.62	0.35 \rightarrow 0.59
Random	30%	0.00 \rightarrow 0.50	0.24 \rightarrow 0.25 \checkmark	0.22 \rightarrow 0.27 \checkmark	0.27 \rightarrow 0.27 \checkmark
Attribution	30%	0.12 \rightarrow 0.65	0.35 \rightarrow 0.69	0.36 \rightarrow 0.56	0.27 \rightarrow 0.55
Group Lasso	30%	0.06 \rightarrow 0.47	0.25 \rightarrow 0.25 \checkmark	0.24 \rightarrow 0.31 \checkmark	0.26 \rightarrow 0.27 \checkmark
Random	63%	0.00 \rightarrow 0.45	0.25 \rightarrow 0.25 \checkmark	0.23 \rightarrow 0.27 \checkmark	0.26 \rightarrow 0.28 \checkmark
Attribution	63%	0.02 \rightarrow 0.52	0.25 \rightarrow 0.30 \checkmark	0.23 \rightarrow 0.34 \checkmark	0.26 \rightarrow 0.32 \checkmark
Group Lasso	63%	0.03 \rightarrow 0.45	0.25 \rightarrow 0.25 \checkmark	0.24 \rightarrow 0.29 \checkmark	0.27 \rightarrow 0.30 \checkmark
Random	80%	0.00 \rightarrow 0.42	0.26 \rightarrow 0.25 \checkmark	0.26 \rightarrow 0.28 \checkmark	0.23 \rightarrow 0.26 \checkmark
Attribution	80%	0.00 \rightarrow 0.47	0.25 \rightarrow 0.25 \checkmark	0.23 \rightarrow 0.27 \checkmark	0.26 \rightarrow 0.28 \checkmark
Group Lasso	80%	0.02 \rightarrow 0.44	0.25 \rightarrow 0.25 \checkmark	0.24 \rightarrow 0.29 \checkmark	0.26 \rightarrow 0.29 \checkmark

6 Related Work

Distillation. Many works have built on the original distillation work of Hinton et al. [26], seeking to transfer intermediate representations [50, 51] and applying these techniques to language models [52, 53], often after pruning [54, 55]. Relevant to our discussion, Turc et al. [56] showed that pretraining the student before distillation can substantially improve results.

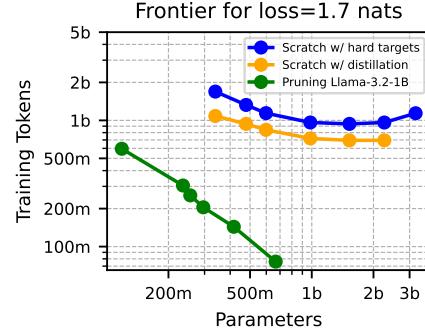


Figure 6: Frontier of network parameters vs. data required to achieve a certain cross-entropy loss on Python documents. For the task of creating a LLM specialized on Python documents, we find that pruning Llama-3.2-1B and then performing recovery training is much more efficient than training LLMs from scratch or distilling LLMs from scratch on the soft targets of Llama-3.2-8B.

Pruning. Even early approaches to pruning used second-order methods for pruning weights [57, 58], whereas our “attribution” scores are first-order. When training vision models, Zhou et al. [25] and Wen et al. [24] used a group lasso penalty like we do, albeit while training on the full data distribution. Sanh et al. [59] proposed “movement pruning” to prune weights during transfer learning. A variety of works have applied structured pruning to LLMs [60, 61], including Xia et al. [62] who develop a method for task-specific pruning. Vural and Erdogdu [63] provide evidence that pruning boosts sample efficiency by directing gradient updates to a sparse set of important directions.

Task Structure and Learning Dynamics. Several works have lately investigated the relationship between task structure and learning dynamics [12–15]. Liu et al. [64] also briefly study a task similar to CMSP to show that hierarchical relationships between tasks cause what they call “domino” learning dynamics. We also highlight the “Skill it!” paper of Chen et al. [65], who empirically investigate hierarchical learning dependencies for real LLM tasks.

Machine Unlearning. Many approaches to unlearning have been proposed [66–70]. Guo et al. [71] study how applying unlearning fine-tuning to different model components affects unlearning success, inspired by a mechanistic understanding of how knowledge is retrieved in LLMs. Lee et al. [72] show that distillation can be applied after unlearning fine-tuning to more robustly remove unwanted capabilities. Highly relevant is the work of Cloud et al. [73], who apply “gradient routing” during training to localize capabilities to different model components, allowing pruning to be used for unlearning.

7 Discussion

Limitations: This work explored some aspects of neural network learning which are relevant to the creation of narrow AI systems. Much work remains to be done in determining how relevant our findings are to real-world networks in various domains. We first emphasize that CMSP is a toy task, an existence proof that curriculum learning effects can be quite strong, but it is still unclear whether these dynamics occur on more natural datasets. For our pruning experiments, we emphasize that our greedy attribution pruning approach is quite simple, and we cannot rule out that more sophisticated pruning strategies would be more successful in preserving some skills while unlearning others [74]. Also, we did not optimally scale hyperparameters in our LLM experiments, so the performance of each of our models in Section 5 is likely imperfect. While we found that pruning a pre-existing Llama model was more efficient than training a model from scratch on Python documents, this could be due to our imperfect training setup, and not because one needs to train on a broad distribution of data to learn Python documents due to curriculum effects.

Conclusion: In this work, we have studied some potential challenges involved in creating narrow AI systems, having to do both with the structure of data and the structures learned internally by neural networks. Underlying this work is a perspective from *mechanistic interpretability*, that neural networks compute a variety of sparse *features* [35], each with a distributed representation [36, 16, 46], and that these features are computed from each other hierarchically in *circuits* [75–77]. First, we found that in order to learn certain complex features, we may have to first train on a broad set of samples which encourage the learning of simpler features. Second, because features are distributed across model components, it is a nontrivial problem to move a set of task-specific features and circuits into a smaller network, especially while *unlearning* unneeded circuits.

While a neural network’s computation across the whole data distribution may be quite complex, we hope that the computation that networks perform on any particular task will be reducible to something less complex. That less complex computation, whatever it is, might be interpretable as a circuit [76, 77], or reduce to a simple program [78], or, as we studied here, could be instantiated in a much smaller network. However, this hope, and the question of whether it is possible to create narrow-purpose versions of today’s models, probes at a basic question about the nature of the intelligence of these models. If their apparent generality indeed results from their having learned a large, diverse set of crystallized, task-specific circuits, then we ought in principle to be able to create competent, specialized versions of these models by just transferring the relevant task-specific circuits. However, if the intelligence of our models, and intelligence more generally, is better understood as resulting from a single unified algorithm, then the basic prospect of creating narrow AI systems that are as strong as truly general ones could be a challenge.

Acknowledgments and Disclosure of Funding

We thank Ziming Liu, Josh Engels, David D. Baek, Tom McGrath, and Jamie Simon for helpful conversations and feedback. E.J.M. was supported by the NSF via the Graduate Research Fellowship Program (Grant No. 2141064) and under Cooperative Agreement PHY-2019786 (IAIFI).

References

- [1] ZZ Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, et al. Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. *arXiv preprint arXiv:2504.21801*, 2025.
- [2] Paul Kassianik, Baturay Saglam, Alexander Chen, Blaine Nelson, Anu Vellore, Massimo Aufiero, Fraser Burch, Dhruv Kedia, Avi Zohary, Sajana Weerawardhena, et al. Llama-3.1-foundationai-securityllm-base-8b technical report. *arXiv preprint arXiv:2504.21039*, 2025.
- [3] Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*, 2024.
- [4] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [5] Fazl Barez, Tingchen Fu, Ameya Prabhu, Stephen Casper, Amartya Sanyal, Adel Bibi, Aidan O’Gara, Robert Kirk, Ben Bucknall, Tim Fist, et al. Open problems in machine unlearning for ai safety. *arXiv preprint arXiv:2501.04952*, 2025.
- [6] Leonard Bereska and Efstratios Gavves. Mechanistic interpretability for ai safety—a review. *arXiv preprint arXiv:2404.14082*, 2024.
- [7] Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, et al. Open problems in mechanistic interpretability. *arXiv preprint arXiv:2501.16496*, 2025.
- [8] Max Tegmark and Steve Omohundro. Provably safe systems: the only path to controllable agi. *arXiv preprint arXiv:2309.01933*, 2023.
- [9] David Dalrymple, Joar Skalse, Yoshua Bengio, Stuart Russell, Max Tegmark, Sanjit Seshia, Steve Omohundro, Christian Szegedy, Ben Goldhaber, Nora Ammann, et al. Towards guaranteed safe ai: A framework for ensuring robust and reliable ai systems. *arXiv preprint arXiv:2405.06624*, 2024.
- [10] Max Tegmark, Søren Mindermann, Vanessa Wilfred, and Wan Sie Lee. The singapore consensus on global ai safety research priorities. Conference report, 2025 Singapore Conference on AI: International Scientific Exchange on AI Safety, Singapore, April 2025. URL <https://file.go.gov.sg/sg-consensus-ai-safety.pdf>.
- [11] Eric Drexler, K. Reframing superintelligence: Comprehensive ai services as general intelligence. Future of Humanity Institute Technical Report, 2019. Available at https://www.fhi.ox.ac.uk/wp-content/uploads/Reframing_Superintelligence_FHI-TR-2019-1.1-1.pdf.
- [12] Rahul Ramesh, Ekdeep Singh Lubana, Mikail Khona, Robert P Dick, and Hidenori Tanaka. Compositional capabilities of autoregressive transformers: A study on synthetic, interpretable tasks. *arXiv preprint arXiv:2311.12997*, 2023.
- [13] Maya Okawa, Ekdeep S Lubana, Robert Dick, and Hidenori Tanaka. Compositional abilities emerge multiplicatively: Exploring diffusion models on a synthetic task. *Advances in Neural Information Processing Systems*, 36:50173–50195, 2023.

[14] Core Francisco Park, Maya Okawa, Andrew Lee, Ekdeep S Lubana, and Hidenori Tanaka. Emergence of hidden capabilities: Exploring learning dynamics in concept space. *Advances in Neural Information Processing Systems*, 37:84698–84729, 2024.

[15] Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2552–2623. PMLR, 2023.

[16] Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.

[17] Paul Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1-2):159–216, 1990.

[18] Geoffrey E Hinton. Learning distributed representations of concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 8, 1986.

[19] Eric Michaud, Ziming Liu, Uzay Girit, and Max Tegmark. The quantization model of neural scaling. *Advances in Neural Information Processing Systems*, 36:28699–28722, 2023.

[20] Neel Nanda. Attribution patching: Activation patching at industrial scale, 2023. URL <https://www.neelnanda.io/mechanistic-interpretability/attribution-patching>.

[21] Aaquib Syed, Can Rager, and Arthur Conmy. Attribution patching outperforms automated circuit discovery. *arXiv preprint arXiv:2310.10348*, 2023.

[22] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.

[23] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(1):49–67, 2006.

[24] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. *Advances in neural information processing systems*, 29, 2016.

[25] Hao Zhou, Jose M Alvarez, and Fatih Porikli. Less is more: Towards compact cnns. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 662–677. Springer, 2016.

[26] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

[27] Boaz Barak, Benjamin Edelman, Surbhi Goel, Sham Kakade, Eran Malach, and Cyril Zhang. Hidden progress in deep learning: Sgd learns parities near the computational limit. *Advances in Neural Information Processing Systems*, 35:21750–21764, 2022.

[28] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.

[29] Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>.

[30] Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. *arXiv preprint arXiv:2301.05217*, 2023.

[31] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.

[32] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

[33] Nayara Fonseca, Seok Hyeong Lee, Chris Mingard, Ard Louis, et al. An exactly solvable model for emergence and scaling laws in the multitask sparse parity problem. *Advances in Neural Information Processing Systems*, 37:39632–39693, 2024.

[34] Itamar Shoshani and Ohad Shamir. Hardness of learning fixed parities with neural networks. *arXiv preprint arXiv:2501.00817*, 2025.

[35] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2017. doi: 10.23915/distill.00007. <https://distill.pub/2017/feature-visualization>.

[36] Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of superposition. *Transformer Circuits Thread*, 2022. https://transformer-circuits.pub/2022/toy_model/index.html.

[37] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. <https://transformer-circuits.pub/2021/framework/index.html>.

[38] Nelson Elhage, Robert Lasenby, and Christopher Olah. Privileged bases in the transformer residual stream. *Transformer Circuits Thread*, page 24, 2023.

[39] Victor Lecomte, Kushal Thaman, Rylan Schaeffer, Naomi Bashkansky, Trevor Chow, and Sanmi Koyejo. What causes polysemanticity? an alternative origin story of mixed selectivity from incidental causes. *arXiv preprint arXiv:2312.03096*, 2023.

[40] Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. Linear algebraic structure of word senses, with applications to polysemy. *Transactions of the Association for Computational Linguistics*, 6:483–495, 2018.

[41] Vikrant Varma, Rohin Shah, Zachary Kenton, János Kramár, and Ramana Kumar. Explaining grokking through circuit efficiency. *arXiv preprint arXiv:2309.02390*, 2023.

[42] Yann LeCun. The mnist database of handwritten digits, 1998. URL <http://yann.lecun.com/exdb/mnist/>.

[43] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

[44] Shuyao Xu, Liu Jiayao, Zhenfeng He, Cheng Peng, and Weidi Xu. The surprising effectiveness of randomness in llm pruning. In *Sparsity in LLMs (SLLM): Deep Dive into Mixture of Experts, Quantization, Hardware, and Inference*.

[45] Yuxiao Li, Eric J Michaud, David D Baek, Joshua Engels, Xiaoqing Sun, and Max Tegmark. The geometry of concepts: Sparse autoencoder feature structure. *Entropy*, 27(4):344, 2025.

[46] Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.

[47] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022.

[48] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

[49] Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D Li, Ann-Kathrin Dombrowski, Shashwat Goel, Long Phan, et al. The wmdp benchmark: Measuring and reducing malicious use with unlearning. *arXiv preprint arXiv:2403.03218*, 2024.

[50] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.

[51] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019.

[52] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.

[53] V Sanh. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.

[54] JS McCarley, Rishav Chakravarti, and Avirup Sil. Structured pruning of a bert-based question answering model. *arXiv preprint arXiv:1910.06360*, 2019.

[55] AI Meta. Llama 3.2: Revolutionizing edge ai and vision with open, customizable models. *Meta AI Blog*. Retrieved March 2025, 2024.

[56] Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*, 2019.

[57] Yann LeCun, John S. Denker, and Sara A. Solla. Optimal brain damage. In David S. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 598–605. Morgan Kaufmann, 1990. URL <https://proceedings.neurips.cc/paper/1989/hash/6c9882bbac1c7093bd25041881277658-Abstract.html>.

[58] Babak Hassibi and David G. Stork. Second order derivatives for network pruning: Optimal brain surgeon. In Stephen J. Hanson, Jack D. Cowan, and C. Lee Giles, editors, *Advances in Neural Information Processing Systems 5*, pages 164–171. Morgan Kaufmann, 1993. URL <https://proceedings.neurips.cc/paper/1992/hash/647-second-order-derivatives-for-network-pruning-optimal-brain-surgeon.html>.

[59] Victor Sanh, Thomas Wolf, and Alexander Rush. Movement pruning: Adaptive sparsity by fine-tuning. *Advances in neural information processing systems*, 33:20378–20389, 2020.

[60] Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, and Danqi Chen. Sheared llama: Accelerating language model pre-training via structured pruning. *arXiv preprint arXiv:2310.06694*, 2023.

[61] Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large language models. *Advances in neural information processing systems*, 36:21702–21720, 2023.

[62] Mengzhou Xia, Zexuan Zhong, and Danqi Chen. Structured pruning learns compact and accurate models. *arXiv preprint arXiv:2204.00408*, 2022.

[63] Nuri Mert Vural and Murat A. Erdogdu. Pruning is optimal for learning sparse features in high-dimensions, 2024. URL <https://arxiv.org/abs/2406.08658>.

[64] Ziming Liu, Yizhou Liu, Eric J Michaud, Jeff Gore, and Max Tegmark. Physics of skill learning. *arXiv preprint arXiv:2501.12391*, 2025.

[65] Mayee Chen, Nicholas Roberts, Kush Bhatia, Jue Wang, Ce Zhang, Frederic Sala, and Christopher Ré. Skill-it! a data-driven skills framework for understanding and training language models. *Advances in Neural Information Processing Systems*, 36:36000–36040, 2023.

[66] Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pages 463–480. IEEE, 2015.

[67] Lucas Bourtoule, Varun Chandrasekaran, Christopher A. Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. *arXiv preprint arXiv:1912.03817*, 2019. doi: 10.48550/arXiv.1912.03817. URL <https://arxiv.org/abs/1912.03817>.

[68] Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *Advances in Neural Information Processing Systems*, 37:105425–105475, 2024.

[69] Jiaao Chen and Diyi Yang. Unlearn what you want to forget: Efficient unlearning for llms. *arXiv preprint arXiv:2310.20150*, 2023.

[70] Kang Gu, Md Rafi Ur Rashid, Najrin Sultana, and Shagufta Mehnaz. Second-order information matters: Revisiting machine unlearning for large language models. *arXiv preprint arXiv:2403.10557*, 2024.

[71] Phillip Guo, Aaquib Syed, Abhay Sheshadri, Aidan Ewart, and Gintare Karolina Dziugaite. Mechanistic unlearning: Robust knowledge unlearning and editing via mechanistic localization, 2024. URL <https://arxiv.org/abs/2410.12949>.

[72] Bruce W. Lee, Addie Foote, Alex Infanger, Leni Shor, Harish Kamath, Jacob Goldman-Wetzler, Bryce Woodworth, Alex Cloud, and Alexander Matt Turner. Distillation robustifies unlearning, 2025. URL <https://arxiv.org/abs/2506.06278>.

[73] Alex Cloud, Jacob Goldman-Wetzler, Evžen Wybitul, Joseph Miller, and Alexander Matt Turner. Gradient routing: Masking gradients to localize computation in neural networks. *arXiv preprint arXiv:2410.04332*, 2024.

[74] Yanyu Li, Pu Zhao, Geng Yuan, Xue Lin, Yanzhi Wang, and Xin Chen. Pruning-as-search: Efficient neural architecture search via channel pruning and structural reparameterization. *arXiv preprint arXiv:2206.01198*, 2022.

[75] Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.001. <https://distill.pub/2020/circuits/zoom-in>.

[76] Samuel Marks, Can Rager, Eric J Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. *arXiv preprint arXiv:2403.19647*, 2024.

[77] Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language model. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/attribution-graphs/biology.html>.

[78] Eric J Michaud, Isaac Liao, Vedang Lad, Ziming Liu, Anish Mudide, Chloe Loughridge, Zifan Carl Guo, Tara Rezaei Kheirkhah, Mateja Vukelić, and Max Tegmark. Opening the ai black box: Distilling machine-learned algorithms into code. *Entropy*, 26(12):1046, 2024.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We clearly state our contributions in our introduction, and these match the contents of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We include a limitations section in Section 7 as well as addressing possible limitations throughout the paper.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Answer: [\[NA\]](#)

Justification: We do not include any theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Yes, we describe our experimental setup in each section. We also have linked to a repository containing the code used to produce all our results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We have linked a github in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We include experimental details in the main text.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[No\]](#)

Justification: In Figure 2, Figure 7, and Figure 10 we directly show results across all seeds instead of plotting the mean and variance. In Figure 4 we currently report means without error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We include compute estimates for our experiments in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: While we are motivated by some broader concerns around societal impact, the work done here is relatively small-scale and exploratory. We discuss positive impacts in the Introduction.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work does not pose such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We properly reference the datasets and models we use.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our work did not involve crowdsourcing nor work with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our work did not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We did not use LLMs in non-standard ways.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Additional results on CMSP

Here we include some additional results on CMSP. We first explore whether network shape influences learning dynamics when learning composite tasks (w/ corresponding atomic tasks in the training dataset) vs. when learning an atomic task on its own. We find that the effect of depth is similar between composite vs. atomic tasks. This suggests that our networks do not necessarily need to be deep in order to take advantage of compositional structure in this task. Note that a previous version of this manuscript suggested the opposite—we had observed that deeper networks learn composite CMSP tasks faster than shallow networks, but it turns out that deeper networks also learn standalone atomic tasks faster than shallow networks. We show loss curves of depth-3 width-128, depth-2 width-361, and depth-2 width-128 networks in Figure 7.

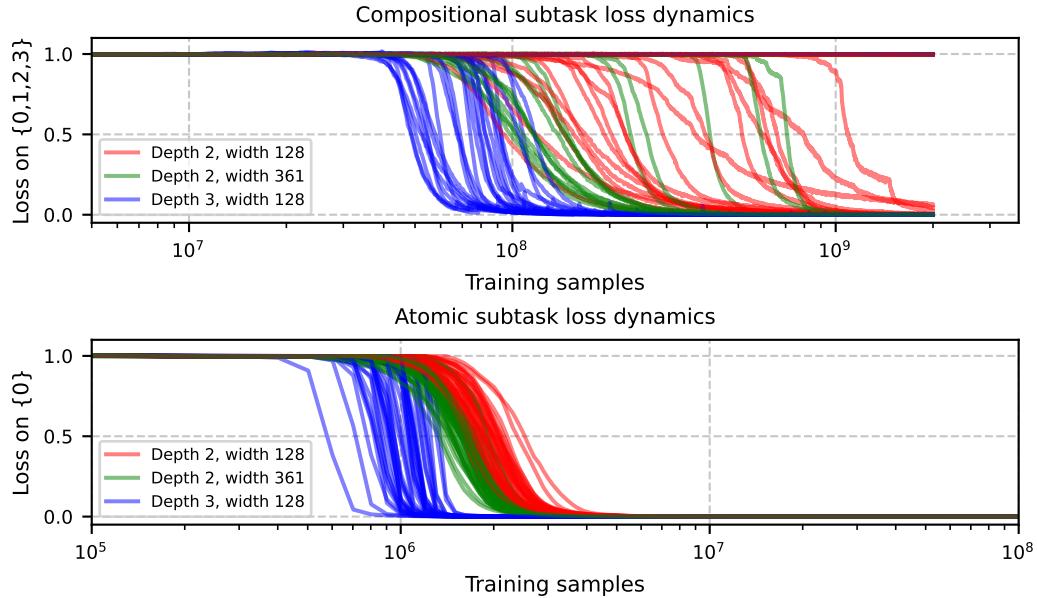


Figure 7: Impact of network shape on CMSP learning. **Top:** Forty runs of training with different seeds on CMSP dataset $\mathcal{D}_{\{0\}} \cup \mathcal{D}_{\{1\}} \cup \mathcal{D}_{\{2\}} \cup \mathcal{D}_{\{3\}} \cup \mathcal{D}_{\{0,1,2,3\}}$ with $m = 4$, $n = 64$, $k = 4$. We find that 27/40 runs converge with two hidden layers versus 19/40 runs with a single hidden layer with width 128 and 16/40 runs with a single hidden layer and width 361. Single-hidden-layer networks are therefore able to learn composite subtasks in CMSP roughly as well as deeper networks. **Bottom:** Forty runs of training on a standalone atomic task $m = 1$, $n = 64$, $k = 4$.

We also show how ablation scores vary across neurons in our unregularized networks. With the setup in Section 3.3, we show ablation scores for each neuron for each subtask in Figure 8 and Figure 9. We find that there are neurons which have high scores across most subtasks, even subtasks in different “skill trees” ($\{0,1,2\}$ versus $\{3,4,5\}$).

In Figure 10 we show sparsity vs accuracy curves from pruning like in Figure 3 (bottom left), across seeds and network sizes. In our unregularized networks, there is a lot of variation between runs in how entangled the different subtasks are.

B Additional results on LLMs

B.1 Additional plots

We also show some additional results with LLMs. In Figure 11, we show how attribution scores correlate with ablation scores in LLMs. We show results for pruning both neurons and residual stream dimensions.

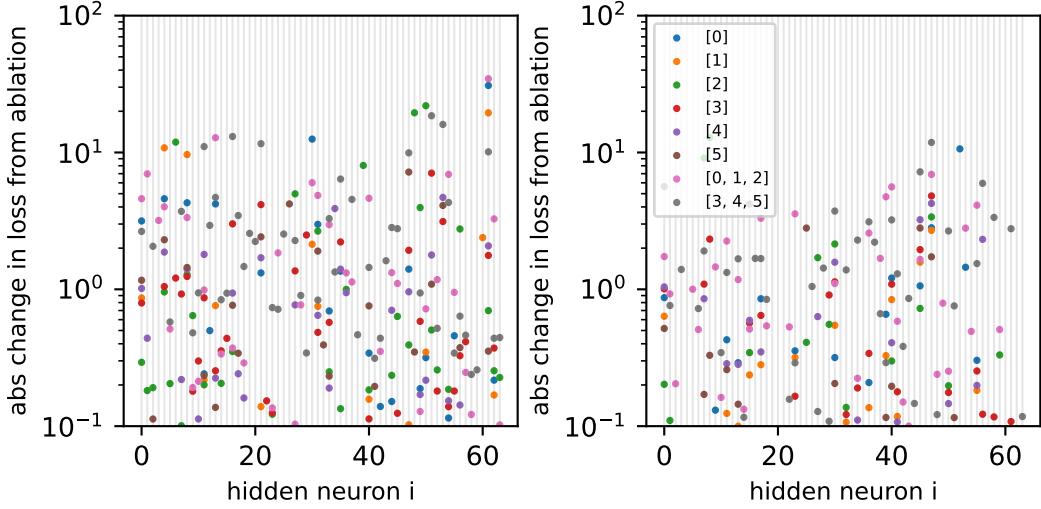


Figure 8: For each hidden neuron in the first hidden layer (**left**) and the second hidden layer (**right**) in a CMSP network, we show the ablation score of that neuron for each subtask. For some neurons, ablation scores are high across multiple subtasks, even from separate “skill trees”.

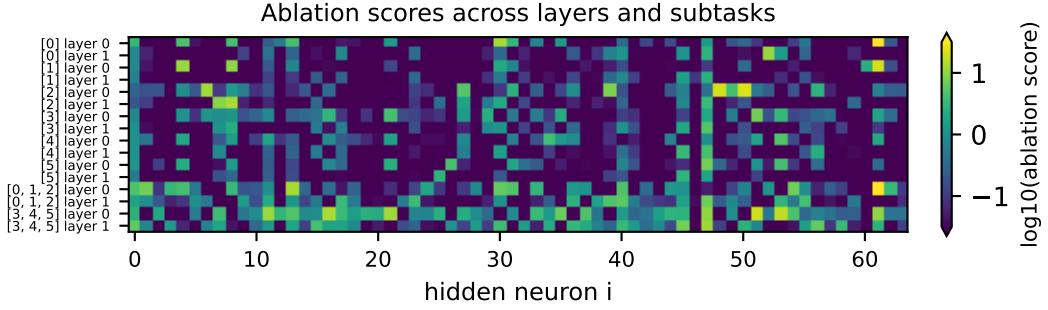


Figure 9: Ablation scores for each subtask and each neuron. This is simply a different way of visualizing the results in Figure 8.

In Figure 12, we show learning curves across runs when training from scratch on Python documents, when performing distillation, and when performing recovery training after pruning both neurons and residual stream dimensions, to varying target sparsities, of Llama-3.2-1B.

In Figure 13, we show recovery curves after pruning neurons and residual stream dimensions from Llama-3.2-1B to varying sparsity levels. We find that pruning random components performs roughly as well as pruning components with the lowest attribution scores.

B.2 Additional experimental details

When training networks from scratch and distilling from scratch (Figure 12 and Figure 6, we train transformers with the Llama 3 architecture [43] of varying size, listed in Table 2).

When we prune Llama-3.2-1B, we prune neurons and residual stream dimensions with the sparsity combinations shown in Table 3.

C Compute estimates

We estimate the compute used for each of our experiments.

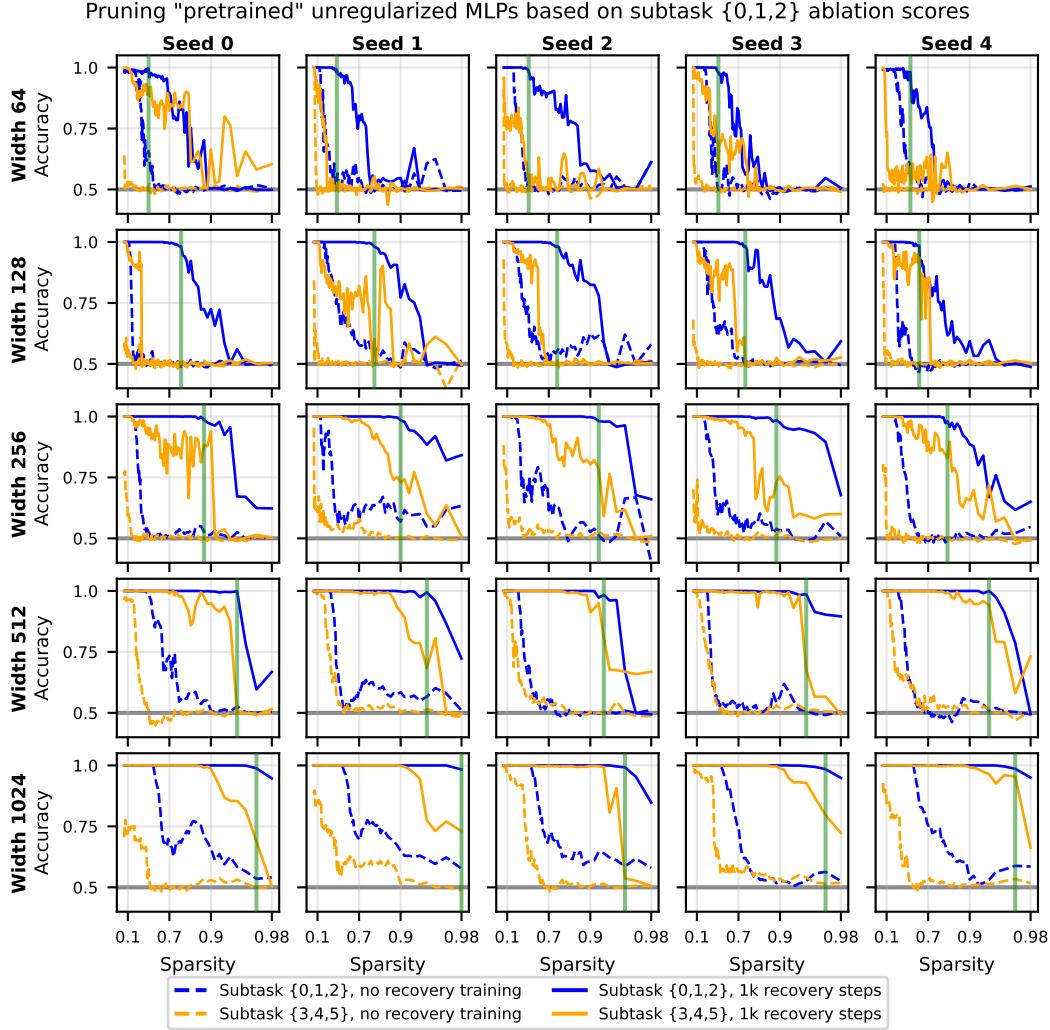


Figure 10: Pruning sparsity vs accuracy curves for pretrained CMSP networks. We train MLPs of varying width and random seed on a CMSP dataset with two skill trees, as in Figure 3. We prune neurons based on ablation score on subtask $\{0,1,2\}$. On some networks, ablating neurons to maximum sparsity while preserving performance on subtask $\{0,1,2\}$ robustly unlearns subtask $\{3,4,5\}$. However, for other networks, the subtasks seem more entangled, and we can recover performance on subtask $\{3,4,5\}$ even after ablating to the maximum extent we can still recover 98% accuracy on subtask $\{0,1,2\}$ (green line).

For the CMSP training experiments shown in Figure 2 and Figure 7, we ran 4 configurations each with 40 different seeds. We ran each job on a GPU, but on a cluster with a variety of different node configurations and GPUs. Jobs generally took between 5-60 minutes to complete, for between 13-160 total hours of GPU time.

For the experiments shown in Figure 3 and Figure 10, each run, across 5 choices of seed and 5 choices of width, trained networks on a CMSP task, ran pruning experiments, performed group-lasso training, and pruned again. Each such run took between 5-120 minutes on our cluster with varying GPU configurations, for a total time of 2-50 hours of GPU time.

For our MNIST experiments shown in Figure 4, we plot a total of 54 points, each of which is averaged over 10 training runs. We estimate that each run took around 2 minutes on our cluster, totaling to around 18 hours of GPU time.

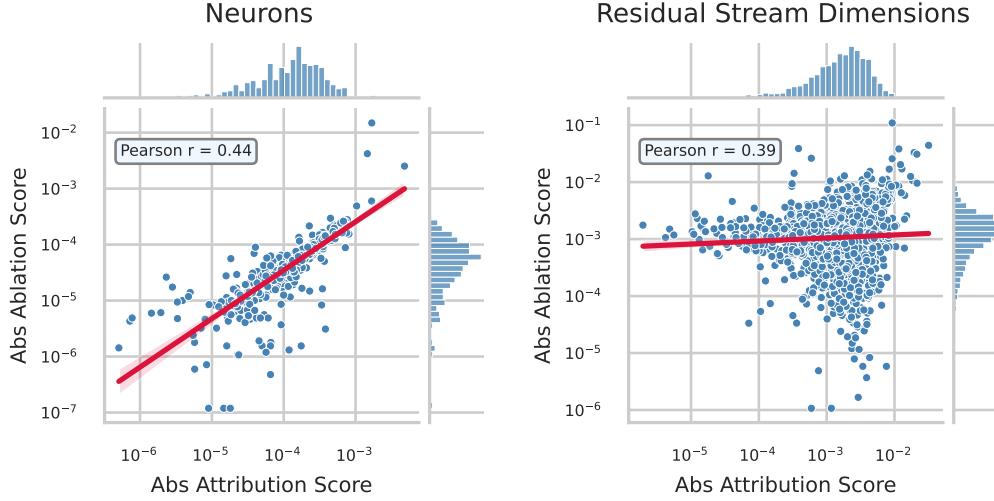


Figure 11: Comparison of ablation vs. attribution scores for Llama-3.2-1B neurons (left) and residual stream dimensions (right), evaluated on a single batch of Python code documents. For each figure, we fully ablate model components and compare the absolute change in loss (Abs Ablation Score, simply called “ablation score” elsewhere in the paper) with the absolute value of the linear estimate of this change computed from model gradients (Abs Attribution Score, called “attribution score” elsewhere). Note that the correlation between the log of these scores is 0.80 for neurons and 0.05 for residual stream components.

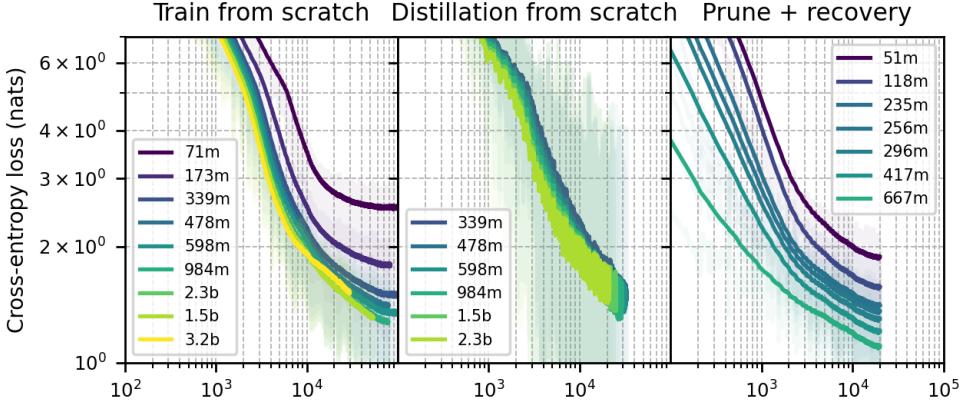


Figure 12: Training curves of LLMs on Python documents when training from scratch (left), training from scratch but with distillation of a larger pretrained LLM (center), and when training pruned pretrained models (right) to recover performance lost during training.

For our LLM experiments, we ran our experiments on A100-80GB nodes, with a single GPU allocated per experiment. When we trained models from scratch on Python documents, we used 9 configurations with job lengths between 1-3 days. For distillation we had 9 configurations with job lengths between 2-3 days, though some jobs failed. For the group lasso training experiments in Figure 5, we show 3 configurations trained with job lengths of 3 days, and we performed recovery training on these models with a job length of 1 day. We estimate that the total time for these jobs was less than 1600 hours of A100 time, though the full set of experiments we attempted in the work that led to this manuscript could be over 5000 hours.

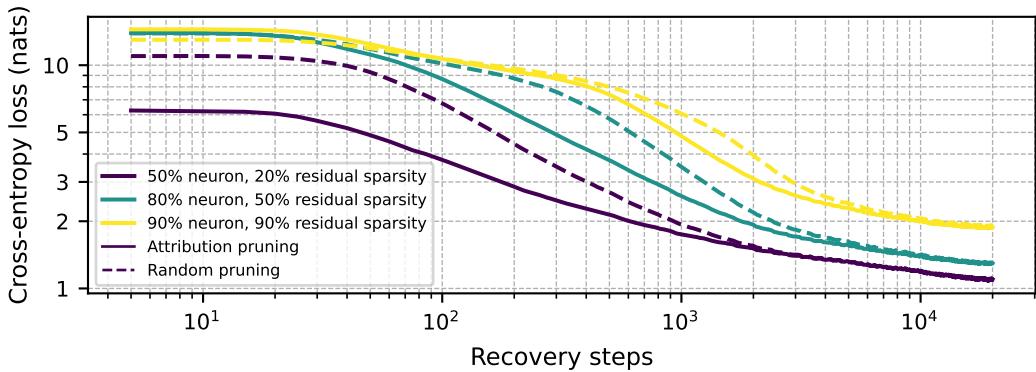


Figure 13: Recovery training curves after pruning neurons and residual stream dimensions, for Llama-3.2-1B on Python code documents. We compare recovery performance when pruning based on attribution scores vs. choosing components randomly. Attribution scores are computed across 1024 documents with a max length of 1024 tokens. We find that pruning with attribution scores is better than pruning random components, however this gap is eventually recovered. For instance, at step 5365 when our run with 50% neuron sparsity and 20% residual sparsity first matches the performance of the original model (≈ 1.3 nats), the performance on randomly-pruned model is almost identical at 1.301 nats.

Table 2: Transformer model configurations explored

Hidden size	#Layers	#Heads	Intermediate size
256	4	4	1,024
512	8	8	2,048
768	12	12	3,072
864	16	16	3,456
1,024	16	16	4,096
1,280	20	20	5,120
1,536	24	24	6,144
1,792	28	28	7,168
2,048	32	32	8,192

Table 3: Neuron and residual sparsity configurations when using attribution pruning on Llama-3.2-1B.

Neuron sparsity	Residual sparsity
0.50	0.50
0.80	0.50
0.90	0.50
0.95	0.50
0.80	0.80
0.90	0.90