

Understanding Language Model Circuits through Knowledge Editing

Huaizhi Ge

Columbia University
hg2590@columbia.edu

Frank Rudzicz

Dalhousie University
frank@dal.ca

Zining Zhu

Stevens Institute of Technology
zzhu41@stevens.edu

Abstract

Recent advances in language model interpretability have identified circuits, critical subnetworks that replicate model behaviors, yet how knowledge is structured within these crucial subnetworks remains opaque. To gain an understanding toward the knowledge in the circuits, we conduct systematic knowledge editing experiments on the circuits of the GPT-2 language model (Radford et al., 2019). Our analysis reveals intriguing patterns in how circuits respond to editing attempts, the extent of knowledge distribution across network components, and the architectural composition of knowledge-bearing circuits. These findings offer insights into the complex relationship between model circuits and knowledge representation, deepening the understanding of how information is organized within language models. Our findings offer novel insights into the “meanings” of the circuits, and introduce directions for further interpretability and safety research of language models.

1 Introduction

The interpretability of large language models (LLMs) has advanced significantly with the exploration of “circuits”—specific subsets of model parameters tailored for distinct tasks. While these circuits, derived from specialized training on text datasets, have proven crucial for replicating model behaviors, a question has not been studied thoroughly: How is knowledge structured within these critical subnetworks? Previous works have approached circuit analysis through ablation studies of designated components (Conmy et al., 2023). Others have attempted to verbalize network components, either using the language model head as a “lens” (nostalgebraist, 2020; Belrose et al., 2023) or an external model (Hernandez et al., 2022; Singh et al., 2023), but these approaches are expensive to scale to the study of the knowledge representation properties of the automatically-extracted circuits.

We propose a novel perspective to understand the knowledge structure within automatically-extracted circuits through systematic knowledge editing experiments. This approach examines circuits as holistic objects rather than analyzing individual components, offering insights into how information is organized and modified within these crucial subnetworks.

Using the GPT-2 base model (Radford et al., 2019), we first investigate knowledge storage patterns within circuits derived from diverse text classification datasets. These circuits were adapted to generation tasks to facilitate knowledge editing experiments. Through knowledge editing, we assess how different circuit components respond to knowledge modifications. Our findings reveal an intriguing “confirmation bias” behavior: Knowledge-intensive circuits demonstrate stronger resistance to editing compared to their complementary parts, suggesting structured information storage patterns.

To better understand knowledge distribution, we explore circuits of varying sizes, from 50% down to 5% of the model’s parameters. This investigation reveals that the ideal knowledge-bearing circuit likely lies between these bounds, indicating that knowledge is neither extremely concentrated nor widely dispersed throughout the network.

Our analysis of cross-dataset circuit overlap provides further insights into knowledge organization. Using mask similarity analysis, we find nuanced patterns of overlap between circuits from different tasks, particularly between linguistic and knowledge-based datasets.

Examining the architectural composition of knowledge-bearing circuits, we make a surprising discovery: LayerNorm components constitute a larger share of the circuits than traditionally emphasized attention and multi-layer perceptron (MLP) layers. This finding provides novel evidence that complements conventional assumptions about knowledge storage in neural networks and high-

lights the potential importance of normalization layers in maintaining network stability and information organization.

These findings collectively enhance our understanding of how neural networks structure and modify information within circuits. By revealing patterns in knowledge distribution, circuit overlap, and architectural composition, our study introduces new directions for interpretability research and potential implications for model safety. Our results suggest that effective manipulation of model behavior may require careful consideration of how knowledge is organized within these crucial subnetworks. All our codes and analysis data is open-sourced at GitHub.

2 Related work

Circuit analysis Circuit extraction tries to find a minimal subnetwork that represents the behavior of the full network (computation graph), where the components can be explainable (Olah et al., 2020; Elhage et al., 2021). This extracted circuit should be faithful (Hanna et al., 2024): if the non-circuit graph edges outside of the circuit are removed, the model’s performance should stay the same. Each component of the circuit (the vertex of the graph) is usually a module (Conmy et al., 2023; Wang et al., 2022) or a causal variable that supports interventions (Vig et al., 2020; Geiger et al., 2021).

While the circuit analysis literature study the module level, alternative granularity levels exist, analyzing on parameter level (Csordás et al., 2020; Bayazit et al., 2023) and sparse autoencoder level (Marks et al., 2024). In this paper, we focus on the parameter level, but our approach can be extended to other granularity levels without loss of generality.

A contemporaneous work, Knowledge Circuits (Yao et al., 2024), also studies the circuit from a knowledge editing point of view. Our works differ in the granularity level of the circuit, the types of knowledge, the tasks studied, and the circuit editing methods.

Model editing Various methods modify knowledge in large language models. Zhu et al. (2020) introduced a task for selective updates of factual knowledge in Transformer models, using fine-tuning as a benchmark. Rank-One Model Editing (ROME) by Meng et al. (2022a) adjusted feed-forward weights to change specific factual associations. MEMIT (Meng et al., 2022b) enabled in-

tegrating multiple memories into a language model. Model Editor Networks with Gradient Decomposition (MEND) by Mitchell et al. (2021) utilized a single targeted input-output pair for rapid, localized changes in a pre-trained model’s behavior. Other influential approaches include targeting specific knowledge neurons (Dai et al., 2021) and using hyper-networks (De Cao et al., 2021). These techniques have proven to be highly effective in editing knowledge within LLMs.

3 Methods

Circuit extraction We adopt a differentiable masking technique implemented in Yu et al. (2024). Similar to the differentiable masking literature (De Cao et al., 2022; Csordás et al., 2020; Bayazit et al., 2023), we set up learning objectives to train a mask parameter $m_\theta \in [0, 1]$ for each model parameter $\theta \in \Theta$, where Θ describes the collection of all trainable parameters in the full model. Unlike De Cao et al. (2022), we use a combination of the following objectives:

1. Faithfulness loss $L_F = \frac{1}{N} \sum_i \sum_j \hat{y}^{(i)} \log \frac{1}{\hat{l}_j^{(i)}}$, which is a cross-entropy loss. This objective encourages $\hat{l}_j^{(i)}$, the logits of the circuit (i.e., the masked model) at class j of sample i , to predict the same results as the original model $\hat{y}^{(i)}$. This loss is averaged across the N training data samples.
2. Sparseness loss $L_S = \frac{1}{|\Theta|} \sum_\theta \sigma(m_\theta)$, where m_θ is the mask for the model parameter θ . This approximately counts the density of the nonzero masks (with a sigmoid function $\sigma(\cdot)$ that smooths the distribution). This encourages the circuit to be sparse.

Following De Cao et al. (2022), we use the hard-concrete distribution of Louizos et al. (2018) to sparsify the masks. Following Yu et al. (2024), we use a straight-through estimator (Bengio et al., 2013) to make the forward pass (apply the mask) trainable. We train until the sparseness reaches a pre-defined threshold.

Circuit-aware model edit In the circuit extraction step, we identified a circuit that is a subset of the model. To study how circuits are affected by the model’s knowledge, we replace the classification head of the model with its language model head, transforming it into a generation model. Thus, we

repurpose the identified circuits from text classification tasks for generation tasks, allowing us to conduct knowledge editing.

This circuit-aware model edit involves altering the true target to a new target. For instance, given the prompt “A cat is a kind of”, we would edit the true target “animal” to the new target “plant”. Throughout this study, we employ fine-tuning techniques to modify the knowledge stored within these circuits.

The circuit-aware model edit is implemented using the fine-tuning method from the EasyEdit repository (Wang et al., 2023), where we modify the knowledge embedded in the model with a cross-entropy loss:

$$L_C = \frac{1}{N} \sum_i \sum_j y^{(i)} \log \frac{1}{p_j^{(i)}}, \quad (1)$$

where N is the number of samples, $y^{(i)}$ is the true label for the i -th sample (0 or 1), and $p_j^{(i)}$ is the predicted probability of the i -th sample being in class j . We use Stochastic Gradient Descent (SGD) to update the parameters:

$$\theta_{t+1} = \theta_t - \eta m_\Theta \nabla_\theta L_C(\theta_t; x^{(i)}, y^{(i)}), \quad (2)$$

where θ_t are the parameters at iteration t , η is the learning rate, and m_Θ is the collection of all binary parameter masks m_θ and is kept constant during the SGD. For parameters involved in the circuit $\mathcal{C} \subseteq \mathcal{P}$, the mask is set to 1; otherwise, it is set to 0. The term $\nabla_\theta L_C(\theta_t; x^{(i)}, y^{(i)})$ represents the gradient of the loss function L_C with respect to the parameters θ_t , computed using a single training example $(x^{(i)}, y^{(i)})$ or a mini-batch of examples.

In this study, we specifically focus on editing the circuits within the model. Following the differential masking literature (Bayazit et al., 2023; Yu et al., 2024), the parameters not involved in these circuits are masked by zero and are maintained at zero throughout the training process by setting their gradients to zero in each step, ensuring they do not affect the model’s behavior.

4 Data

4.1 Model edit data

We use HIERARCHY DATA (Ge et al., 2024) to act as the model edit data. This dataset encompasses a series of both challenging incorrect facts, represented as (s, r, o^*) , and their corresponding accurate facts, denoted as (s, r, o) . Here, s signifies the

subject and o the object, both selected from the hierarchy chains. The relation r consistently adopts the hypernym schema, emphasizing hierarchical connections. An example of this is “A British Shorthair is a kind of cat”. We modify the objects of these facts, generating altered facts “A British Shorthair is a kind of dog” to test the efficacy of circuit edits. The HIERARCHY DATA dataset includes 198 facts targeted for editing analysis. This structured approach facilitates explorations into the role of hierarchical relations in the adaptability and accuracy of language model editing processes.

4.2 Text classification data

We include a range of datasets spanning hierarchical concept knowledge, linguistic knowledge, and those involving model’s safety. Each dataset is reformatted into a classification task that facilitates circuit extraction.

Hierarchical data Here we re-format HIERARCHY DATA(h) into a classification dataset. For instance, we categorized the statement “A cat is a kind of animal” with a label of 1, indicating a true assertion, whereas the statement “A cat is a kind of plant” was labeled as 0, denoting a false assertion.

Linguistic knowledge We include eight datasets from BLiMP (Warstadt et al., 2020). These cover both the syntax and semantic aspects, allowing us to examine if the edits affect the linguistic knowledge of the language models. The datasets are:

- determiner noun agreement 1 (dnal)
- determiner noun agreement 2 (dna2)
- determiner noun agreement irregular 1 (dnai1)
- determiner noun agreement irregular 2 (dnai2)
- determiner noun agreement with adjective 1 (dnawa1)
- determiner noun agreement with adj 2 (dnawa2)
- determiner noun agreement with adj irregular 1 (dnawai1)
- determiner noun agreement with adj irregular 2 (dnawai2)

Safety behavior A crucial goal for the interpretability analysis of models is to make the models safer. To this end, we use five datasets from Perez et al. (2022) that describe the safety-related behavior of language models in the real world, including:

- risk averse (ra)

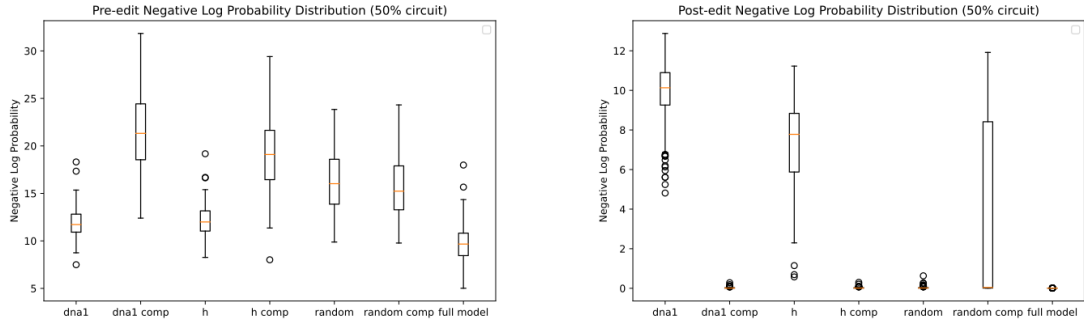


Figure 1: Negative log probability distribution of the 50% circuits. Left: distribution of the true targets pre-edit. Right: distribution of the new targets post-edit.

- extraversion (e)
- desire for being rated helpful, harmless & honest over actually being helpful, harmless & honest (hhh)
- okay with using many resources (umr)
- desire for advancing technology to achieve goals (tech)

5 Validating the circuit extraction

To validate the accuracy and correctness of the extracted circuits, we evaluate their performance across specific classification tasks. In particular, we focus on two types of circuits: those containing 5% parameters and those containing 50% parameters.

Table 1 presents the task accuracy achieved by these identified circuits. The circuits were evaluated based on their ability to reproduce task outcomes. High accuracy scores indicate that the circuits not only preserve the task-relevant information but also reflect the underlying computational mechanisms of the original model. The results show that both the 5% and 50% circuit subsets achieve high levels of accuracy, suggesting that the circuits have been correctly identified.

6 How knowledgeable are the LLM circuits?

6.1 Experiments

To investigate the amount of knowledge in the LLM circuits, we edit the knowledge within a specific circuit and compare its performance with that of its complementary circuit. Typically, individual circuits represent only a small percentage of the model’s total parameters. However, for a clearer demonstration of the amount of knowledge in the circuits, we analyze the 50% circuits, ensuring that

these had an equivalent number of parameters as their complementary circuits.

In addition, we also refer to the remainder part of the model’s parameters “the complementary circuit”. For example, the complementary circuit of the dna1 task is referred to as “dna1 complementary” (dna1 comp for short).

6.2 The circuit vs complementary circuit

Even before applying any edits, we want to understand how the circuit and the complementary circuit encode entity-related knowledge, so we query the pre-edit negative log probabilities of the true targets. A lower values indicate a stronger grasp of the knowledge.

Figure 1 (left) displays the results. The extracted circuits demonstrate significant lower ($p < 0.05$) negative log probabilities of the HIERARCHYDATA entities compared to their complementary circuits, indicating a superior understanding of the knowledge. Note that none of the circuits outperforms the full model, showing that the complementary circuits still capture some factual knowledge.

A noteworthy observation is that this effect generalizes across both the h circuit (which is derived from the same data as the HIERARCHYDATA entities) and the dna1 circuit (which is tasked with a distinct, syntax-focused dataset). This finding suggests that the model’s circuits for solving the syntax tasks may be relevant to the model’s mechanism for storing knowledge, thus providing a potential explanation for the “spill-over” effects (Sahak et al., 2023) — i.e., the effects of some datasets on the models may “spill over” to other apparently irrelevant datasets. We will study the overlap of circuits between different tasks in Section 8.

As a baseline, the probability of a random circuit is not significantly different from that of its

	h circuit	dna1 circuit	dna2 circuit	ra circuit	e circuit
5% circuit	0.84	0.91	0.87	0.94	0.97
50% circuit	0.88	0.89	0.87	0.94	0.87

Table 1: Task accuracy of identified circuits

complementary circuit, confirming that the aforementioned “circuit vs complementary” differences do not arise from randomness.

6.3 The complementary circuit is more susceptible to the model editing

After applying the circuit-aware knowledge editing, we query the negative log probabilities of the HIERARCHYDATA entities again, and the outcomes are presented in Figure 1 (right).

Our metric for assessing edit performance is the negative log probability of the new target, where lower values indicate better editing. The figure displays results for circuits extracted using the “determiner noun agreement 1” (dna1) dataset of BLiMP, circuits extracted using hierarchy data (h) along with their complementary circuits, a randomly extracted circuit alongside its complementary circuit, and the full model circuit. Further details about other datasets are provided in Appendix A. Notably, the complementary 50% circuit consistently outperforms the 50% circuit in terms of editing for both the dna1 and h circuits. Furthermore, the performance of the random circuit is similar to that of its complementary counterpart. These indicate that the complementary circuits, which are less task-relevant, are more susceptible to knowledge edits. Yet, the full model demonstrates the best editing performance, indicating that the circuit might still synergize with the complementary networks upon the injection of new knowledge.

Similar to the pre-edit observations, the post-edit discrepancies between the circuits and their complements can be generalized to a linguistic task (dna1). This suggests that task-specific circuits might indicate the knowledge storage locations. In other words, altering the knowledge from these knowledge-rich circuits is more difficult than altering the knowledge from their complements.

Interestingly, humans have similar “confirmation bias” behavior, tending to reject hypotheses that do not agree with prior assumptions (Wason, 1960). Our findings here can also provide a novel perspective about neural network plasticity (Lyle et al., 2023), but a further in-depth inquiry into the

nature of circuit-specific plasticity is left to future works.

7 How large are “ideal circuits”?

Suppose that there exists a circuit that stores all knowledge needed to perform the task — let’s use the term “ideal circuit” to describe it — how large would it be?

In our previous experiments, we conducted experiments on circuits that comprised 50% of the total model parameters. To investigate potential disparities, we reduced our analysis to include circuits that accounted for 5% of the parameters, using HIERARCHYDATA.

Figure 2 (left) presents the negative log probability scores of the true targets before model edits, and Figure 2 (right) illustrates the outcomes of editing the 5% vs. the 50% circuits, along with their corresponding complementary circuits. We continue to employ the negative log probability of the new targets to quantify the edit performance. Detailed figures of 5% circuits derived from other datasets are presented in Appendix B.

Figure 2 shows that both the 5% and the 50% circuits are harder to edit than their corresponding complementary circuits. Notably, the 5% circuit displays a higher level of such “confirmation bias” than the 50% circuit. This suggests that the 5% circuit may be closer to an ideal circuit, than the 50% circuit.

Now, if an ideal circuit exists, would its size be smaller than 5% or between 5% and 50%?

As the pre-edit logprob distribution plot in Figure 2 (left) shows, the 5% circuit captures the knowledge the best, followed by the 50% circuit (perhaps because the 50% circuit may contain non-knowledgeable “noise” parameters that mask out the circuit’s ability to demonstrate knowledge). Also shown in this figure, the complementary of 5% circuit is more knowledgeable than the complementary of 50% circuit, indicating that the 5% complementary circuit likely possesses certain key parameters directly associated with knowledge retention, whereas the 50% complementary circuit

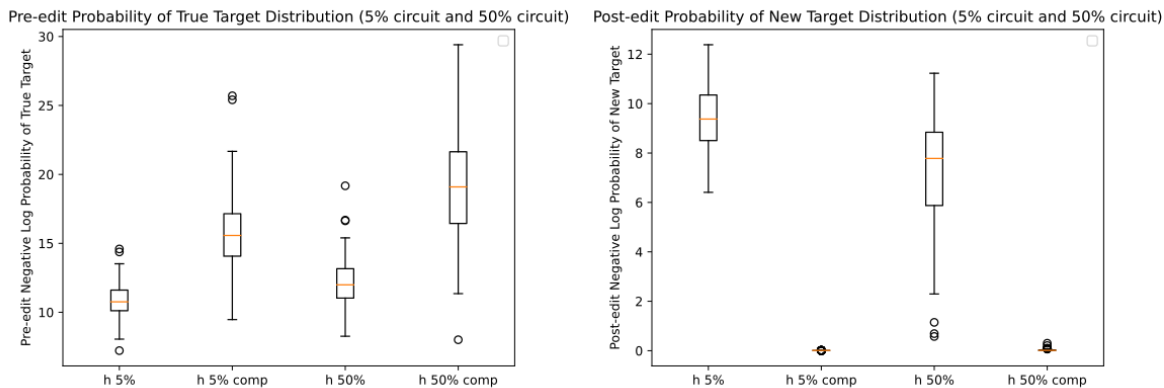


Figure 2: Pre-edit (left) and post-edit (right) log-probability distribution of 5% circuit and 50% circuit

may include a surplus of less critical parameters. In other words, an ideal circuit, if it exists, would incorporate more than 5% but less than 50% of parameters.

Table 2 provides additional evidence toward the sizes of the idealized circuits. We replicate the experiments, this time extracting 15%, 25%, 35% parameters circuits while using the same hyperparameters. The table below shows the median post-edit negative log probabilities of the new target, combined with the results from the 50% and 5% circuits (higher values indicate worse performance). As expected, the edit performance of the extracted circuits declines as the circuit size decreases. If an ideal circuit exists that encapsulates all relevant knowledge in the model, it would likely have the worst edit performance due to strong “confirmation bias”. Therefore, we hypothesize that the 5% circuit may be closer to this “ideal circuit” than the 50% circuit. However, the exact size of this idealized circuit remains to be further investigated.

Recent papers on circuits detected highly sparse circuits that usually contain less than 1% of the total parameters (Bayazit et al., 2023) or less than 3% of total edges (Hanna et al., 2024) for replicating the model behavior. We hypothesize that while the circuits less than 1% explain the functions, the “ideal circuits” need more parameters that correspond to the knowledge used in the tasks. Also, considering that the knowledge localization might differ from the model editing (Hase et al., 2024), the roles of distinct parameters in the model would be an intriguing path to explore further.

8 Do these circuits overlap?

The datasets used for circuit extraction in the previous sections showcase considerable diversity, yet they yield surprisingly consistent results. This consistency raises an interesting question: do these circuits significantly overlap?

To create a baseline, we randomly select a 5% parameter circuit mask and calculate the overlap percentage of all parameters within the circuit. As shown in Table 3, the percentages constantly appear around 5% and only deviate by errors that are two magnitudes smaller. Mathematically, the expectation of a circuit with X% parameter overlaps with a random X% circuit is X%.

Figure 3 features a heatmap that depicts the mask similarity across different tasks. Each entry in the heatmap represents the degree of overlap between 5% circuits from two datasets, measured relative to the base dataset listed in the row. For example, a value of 0.137 at the intersection of the second row and the first column suggests that the circuit derived from the dna1 dataset overlaps with the h dataset circuit by 13.7% of the dna1 circuit.

The degrees of overlap between the circuits vary across different dataset comparisons — between 13% and 18% — which are at least twice the baseline of 5%. These indicate non-negligible similarities. Interestingly, datasets that appear closely related, such as dna1 and dna2, sometimes display significant differences when viewed through the lens of the model’s mechanism. Notably, the ra (risk-averse) circuit shows greater overlaps with other circuits, suggesting that the risk-averse dataset may include terms that are broadly relevant across multiple contexts. Conversely, the e (extraversion) circuit demonstrates fewer overlaps, im-

	50% circuit	35% circuit	25% circuit	15% circuit	5% circuit
Circuit	7.78	8.10	8.28	9.30	9.37
Complementary circuit	1.26e-2	9.99e-2	9.97e-2	9.85e-2	9.81e-2

Table 2: Median post-edit negative log probabilities of the new target across circuits with different sizes

Task	dna1	dna2	h	dnai1	dnawa1	dnawai1	e	ra	hhh
Difference from 0.05	2.6e-4	2.4e-4	1.3e-4	0.7e-4	2.1e-4	2.3e-4	0.1e-4	-0.3e-4	0.7e-4

Table 3: Percentage of parameter overlap between each 5% circuit and a randomly selected circuit




Figure 3: Mask similarity across tasks heatmap. The overlaps are neither too large nor too small.

plying that the extraversion dataset likely features specialized terms that do not commonly occur in other datasets.

Similarly, the 50% circuit of three datasets overlap with each other at levels significantly higher than the 50% expectation, as shown in Table 4.

Our finding echoes those of Merullo et al. (2023) that the circuits of different tasks may overlap with each other. Our analysis also underscores nuanced relationships between different datasets and highlights how specific dataset characteristics might influence the generalizability of the circuits derived from them.

9 What Transformer modules are in the circuits?

In this section, we identify which types of layers are particularly influential within the extracted circuit. Table 5 presents a detailed breakdown of the proportion of total parameters that each type of layer accounts for within the circuit. The data re-

veal that LayerNorm constitutes a large percentage of the circuit’s parameters, especially noticeable within the 5% circuit subset. In contrast, the attentions and multi-layer perceptrons (MLPs) account for a smaller, yet comparable, share of the circuits.

This distribution of parameters is intriguing because it suggests that, while LayerNorm plays a crucial structural role, the computational responsibilities might be more evenly distributed or perhaps differently focused in attention and MLP layers. Notably, some studies found that the MLPs are critical repositories of knowledge within neural networks (Geva et al., 2020, 2022), and the attentions are critical for the mechanisms related to factual recall (Geva et al., 2023). Recently, Katz and Belinkov (2023) found that LayerNorm is related to the introduction of new tokens. There are also some recent debates about the exact localization of the knowledge (Hase et al., 2024; Niu et al., 2024). Complementing prior works, our observations show that LayerNorm plays a critical role in the circuits related to the syntax and hierarchical knowledge entity tasks.

The significant allocation of parameters to LayerNorm in the circuits raises questions about their specific function and impact on the overall network efficacy. Future analysis of LayerNorm could further advance our understanding.

10 Discussion

Safety-related model behavior is relevant to the factual knowledge stored in the LLMs The relationship between safety-related model behavior and the factual knowledge stored in large language models (LLMs) is intricate and significant. During the circuit extraction phase, we utilized model behavior datasets, such as those indicating ‘risk aversion’ and ‘extraversion’, in Model-Written Evaluations. These behavior-related circuits appear to

	h circuit	dna1 circuit	ra circuit
h circuit	1	0.61	0.58
dna1 circuit	0.61	1	0.59
ra circuit	0.58	0.59	1

Table 4: Percentage of parameter overlap between 50% parameter circuits

	dna1 5% circuit	dna1 50% circuit	h 5% circuit	h 50% circuit
attn	0.04	0.30	0.04	0.47
mlp	0.05	0.30	0.05	0.50
ln	0.37	0.63	0.37	0.62

Table 5: Proportion of total parameters accounted for by each layer type within the circuit

possess a substantial amount of factual knowledge. Understanding LLM safety fundamentally involves comprehending model behavior. To achieve this, it is essential to understand the knowledge embedded within the model. However, the precise nature of the knowledge or concepts required for models to determine their behaviors remains an open question. Additionally, the mechanisms by which models use these concepts to guide their behavior are not fully understood. These areas represent critical open problems in the field.

Remove biases without compromising the model’s performance

One intriguing finding is that complementary circuits are easier to edit than the primary behavior-related circuits we identified. This suggests the possibility of performing editing tasks on complementary circuits without altering the primary circuits, thereby preserving the model’s behavior. If validated, this approach could offer a method to remove biases through editing tasks without compromising the model’s performance by using circuit identification.

Implications for model training The understanding of the information structures in the circuits can lead to fine-tuning technologies. The neural network circuit is a particularly appropriate medium because the sparsity can lead to high efficiency, and the automatic circuit discovery algorithms allow the search of circuits that sacrifice model performance to the fewest possible extent. We plan to continue developing the training algorithms in future works.

Recommendations for safer models We propose several recommendations for future researchers aiming to enhance model safety. First,

further investigation into the role and characteristics of complementary circuits can provide insights into how they interact with primary circuits and affect overall model behavior. Second, investigating the possibility of ideal circuits that can perform well for every specific task. Third, exploring other techniques like ROME (Meng et al., 2022a) to edit complementary circuits without impacting primary circuits may help mitigate biases and improve model fairness. Fourth, creating more comprehensive datasets that capture a wider range of behaviors can enhance our understanding of the relationships between behaviors and the knowledge stored in LLMs. Along these paths, future research can contribute to creating safer and more reliable language models.

11 Conclusion

In this paper, we combine two model analysis techniques: knowledge edits and circuit analysis, to inspect the model’s knowledge structure. Our comprehensive experiments reveal novel properties in the automatically extracted circuits, including the resistance to knowledge edit, which we compare to the “confirmation bias”. Our experiments reveal distribution patterns of the knowledge within circuits, across the circuits, and the roles of the knowledge in safety behavior tasks. These findings support future explorations into model behaviors, bias removal, and interpretability. By addressing the open problems and recommendations outlined in our discussion, future research can contribute to creating safer, fairer, and more reliable language models.

12 Limitation

We used only the GPT-2 model. Its structure (Transformer decoder) is the backbone of the current LLMs, and the modules are widely used. We only investigated fine-tuning model edit, and alternative model editing methods can be explored. We investigated a handful of text classification problems including linguistic tasks and safety behavior tasks, but there can always be more.

13 Acknowledgements

A previous version of this document contained a hidden prompt entered by Z Zhu without knowledge of – or consent by – his co-authors. This version does not contain the prompt.

References

- Deniz Bayazit, Negar Foroutan, Zeming Chen, Gail Weiss, and Antoine Bosselut. 2023. [Discovering Knowledge-Critical Subnetworks in Pretrained Language Models](#). *Preprint*, arxiv:2310.03084.
- Nora Belrose, Zach Furman, Logan Smith, Danny Hahawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. [Eliciting Latent Predictions from Transformers with the Tuned Lens](#). *Preprint*, arxiv:2303.08112.
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Arthur Conmy, Augustine Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. 2023. Towards automated circuit discovery for mechanistic interpretability. *Advances in Neural Information Processing Systems*, 36:16318–16352.
- Róbert Csordás, Sjoerd van Steenkiste, and Jürgen Schmidhuber. 2020. Are neural nets modular? inspecting functional modularity through differentiable weight masks. *arXiv preprint arXiv:2010.02066*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2021. Knowledge neurons in pretrained transformers. *arXiv preprint arXiv:2104.08696*.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. Editing factual knowledge in language models. *arXiv preprint arXiv:2104.08164*.
- Nicola De Cao, Leon Schmid, Dieuwke Hupkes, and Ivan Titov. 2022. [Sparse interventions in language models with differentiable masking](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 16–27, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1:1.
- Huaizhi Ge, Frank Rudzicz, and Zining Zhu. 2024. [How Well Can Knowledge Edit Methods Edit Perplexing Knowledge?](#) *arXiv preprint arXiv: 2406.17253*.
- Atticus Geiger, Hanson Lu, Thomas Icard, and Christopher Potts. 2021. Causal abstractions of neural networks. *Advances in Neural Information Processing Systems*, 34:9574–9586.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. Dissecting recall of factual associations in auto-regressive language models. *arXiv preprint arXiv:2304.14767*.
- Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. 2022. [Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2020. Transformer feed-forward layers are key-value memories. *arXiv preprint arXiv:2012.14913*.
- Michael Hanna, Sandro Pezzelle, and Yonatan Belinkov. 2024. Have faith in faithfulness: Going beyond circuit overlap when finding model mechanisms. *arXiv preprint arXiv:2403.17806*.
- Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. 2024. Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models. *Advances in Neural Information Processing Systems*, 36.
- Evan Hernandez, Sarah Schwettmann, David Bau, Teona Bagashvili, Antonio Torralba, and Jacob Andreas. 2022. Natural Language Descriptions of Deep Visual Features. In *ICLR*.
- Shahar Katz and Yonatan Belinkov. 2023. [VISIT: Visualizing and interpreting the semantic information flow of transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14094–14113, Singapore. Association for Computational Linguistics.
- Christos Louizos, Max Welling, and Diederik P Kingma. 2018. Learning sparse neural networks through l₀ regularization. In *International Conference on Learning Representations*.
- Clare Lyle, Zeyu Zheng, Evgenii Nikishin, Bernardo Avila Pires, Razvan Pascanu, and Will Dabney. 2023. [Understanding plasticity in neural networks](#). *Preprint*, arxiv:2303.01486.

- Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. 2024. [Sparse Feature Circuits: Discovering and Editing Interpretable Causal Graphs in Language Models](#). *arXiv preprint arXiv:2211.00593*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2022b. Mass editing memory in a transformer. *arXiv preprint arXiv:2210.07229*.
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. 2023. [Circuit Component Reuse Across Tasks in Transformer Language Models](#). *Preprint*, arxiv:2310.08744.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2021. Fast model editing at scale. *arXiv preprint arXiv:2110.11309*.
- Jingcheng Niu, Andrew Liu, Zining Zhu, and Gerald Penn. 2024. What does the knowledge neuron thesis have to do with knowledge? *arXiv preprint arXiv:2405.02421*.
- nostalgebraist. 2020. [Interpreting GPT: The logit lens](#).
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. 2020. Zoom in: An introduction to circuits. *Distill*, 5(3):e00024–001.
- Ethan Perez, Sam Ringer, Kamilė Lukošiušė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, et al. 2022. Discovering language model behaviors with model-written evaluations. *arXiv preprint arXiv:2212.09251*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Esmat Sahak, Zining Zhu, and Frank Rudzicz. 2023. [A state-vector framework for dataset effects](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15231–15245, Singapore. Association for Computational Linguistics.
- Chandan Singh, Aliyah R. Hsu, Richard Antonello, Shailee Jain, Alexander G. Huth, Bin Yu, and Jianfeng Gao. 2023. [Explaining black box text modules in natural language with language models](#).
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. *Advances in neural information processing systems*, 33:12388–12401.
- Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. *arXiv preprint arXiv:2211.00593*.
- Peng Wang, Ningyu Zhang, Xin Xie, Yunzhi Yao, Bozhong Tian, Mengru Wang, Zekun Xi, Siyuan Cheng, Kangwei Liu, Guozhou Zheng, et al. 2023. Easyedit: An easy-to-use knowledge editing framework for large language models. *arXiv preprint arXiv:2308.07269*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [Blimp: The benchmark of linguistic minimal pairs for english](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Peter C Wason. 1960. On the failure to eliminate hypotheses in a conceptual task. *Quarterly journal of experimental psychology*, 12(3):129–140.
- Yunzhi Yao, Ningyu Zhang, Zekun Xi, Mengru Wang, Ziwen Xu, Shumin Deng, and Huajun Chen. 2024. [Knowledge Circuits in Pretrained Transformers](#). *arXiv preprint*.
- Lei Yu, Jingcheng Niu, Zining Zhu, and Gerald Penn. 2024. [Functional Faithfulness in the Wild: Circuit Discovery with Differentiable Computation Graph Pruning](#). *arXiv preprint arXiv: 2407.03779*.
- Chen Zhu, Ankit Singh Rawat, Manzil Zaheer, Srinadh Bhojanapalli, Daliang Li, Felix Yu, and Sanjiv Kumar. 2020. Modifying memories in transformer models. *arXiv preprint arXiv:2012.00363*.

A 50% circuits probability distribution

This section presents the probability distributions for circuits discovered using different datasets, each representing 50% of the model’s parameters. The datasets used include BLiMP and Model-Written Evaluations. For this experiment, we utilized a GPU with 48GB of memory to perform the computations, with each dataset requiring less than one hour to process. The figures illustrate the post-edit negative log probability distributions of the new target distributions and the pre-edit probability distribution of the true target, providing insights into the impact of circuit-aware model editing.

Figures 4 and 5 display the post-edit probability of the new target and the pre-edit probability of the true target for the BLiMP dataset, respectively. Similarly, Figures 6 and 7 show these distributions for the Model-Written Evaluations dataset.

Overall, these visualizations provide a comprehensive view of the 50% circuits concerning knowledge editing.




Figure 4: Post-edit Negative Log Probability of New Target Distribution (50% circuit) (BLiMP).




Figure 5: Pre-edit Negative Log Probability of True Target Distribution (50% circuit) (BLiMP).

B 5% circuits probability distribution

This section presents the probability distributions for circuits discovered using different datasets, which comprise 5% of the model’s parameters. The datasets used include BLiMP and Model-Written Evaluations. For this experiment, we also utilized a GPU with 64GB of memory to perform the computations, with each dataset requiring less than one hour to process. The figures illustrate the post-edit




Figure 6: Post-edit Negative Log Probability of New Target Distribution (50% circuit) (Model-Written Evaluations).




Figure 7: Pre-edit Negative Log Probability of True Target Distribution (50% circuit) (Model-Written Evaluations).

negative log probability distributions of the new target distributions and the pre-edit probability distribution of the true target, providing insights into the impact of circuit-aware model editing.

Figures 8 and 9 show the post-edit probability of the new target and the pre-edit probability of the true target for the BLiMP dataset, respectively. Figures 10 and 11 display the post-edit probability of the new target and the pre-edit probability of the true target for the Model-Written Evaluations dataset, respectively.

These visualizations offer a detailed perspective on the performance of the 5% circuits in the context of knowledge editing.




Figure 8: Post-edit Negative Log Probability of New Target Distribution (5% circuit) (BLiMP).




Figure 9: Pre-edit Negative Log Probability of True Target Distribution (5% circuit) (BLiMP).




Figure 10: Post-edit Negative Log Probability of New Target Distribution (5% circuit) (Model-Written Evaluations).




Figure 11: Pre-edit Negative Log Probability of True Target Distribution (5% circuit) (Model-Written Evaluations).