

Repetition Neurons: How Do Language Models Produce Repetitions?

Tatsuya Hiraoka Kentaro Inui

Mohamed bin Zayed University of Artificial Intelligence (MBZUAI)
RIKEN

{tatsuya.hiraoka, kentaro.inui}@mbzuai.ac.ae

Abstract

This paper introduces **repetition neurons**, regarded as “skill neurons” responsible for the repetition problem in text generation tasks. These neurons are progressively activated more strongly as repetition continues, indicating that they perceive repetition as a task to copy the previous context repeatedly, similar to in-context learning. We identify these repetition neurons by comparing activation values before and after the onset of repetition in texts generated by recent pre-trained language models. We analyze the repetition neurons in three English and one Japanese pre-trained language models and observe similar patterns across them.

1 Introduction

While text generation with LLMs such as GPT-3 (Brown et al., 2020) has been actively studied, the issue of repetition remains a fundamental challenge (Li et al., 2023a; Ivgi et al., 2024). Specifically, repetition is particularly problematic under greedy generation, which is often used when reproducibility must be guaranteed (Song et al., 2024).

Many researchers have tackled this problem by analyzing repetition (Fu et al., 2021; Xu et al., 2022) and developing techniques to mitigate repetitive outputs (Keskar et al., 2019; Shirai et al., 2021; Zhu et al., 2023; Li et al., 2023a). Some works specifically focus on attention heads, such as induction heads, framing repetition as a key mechanism for in-context learning (Olsson et al., 2022; Bansal et al., 2023; Crosbie and Shutova, 2024). However, the internal mechanisms of generative models that produce repetitive outputs remain insufficiently explored (Vaidya et al., 2023; Wang et al., 2024).

We focus on the neurons of Transformer language models (Vaswani et al., 2017; Geva et al., 2021; Dai et al., 2022; Chen et al., 2024) that detect repetition in inputs and trigger repetitive outputs in text generation. We refer to these neurons as “**repetition neurons**” following Wang et al. (2024).

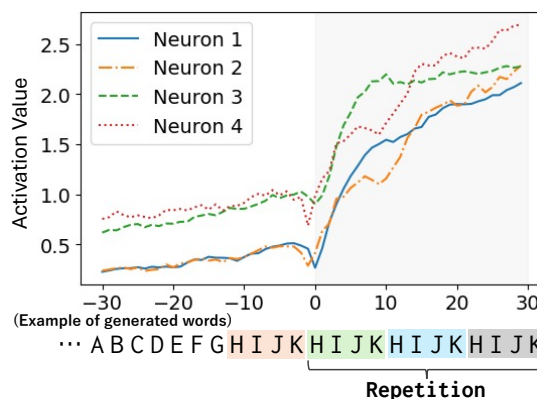


Figure 1: Activation values of top four repetition neurons for 30 tokens before and after repetition (Gemma-2B, averaged value over 1,000 texts). **Repetition neurons are strongly activated in the repetition range.**

We hypothesize that the repetition neuron is a “skill neuron” (Radford et al., 2017; Wang et al., 2022) that prompts the model to generate repetition as a task of copying the previous context, akin to “task vectors” (Hendel et al., 2023) found in in-context learning (Brown et al., 2020; Yan et al., 2024).

We propose a method to identify repetition neurons by comparing activation values in the input ranges before and after the onset of repetition (§3.1). As shown in Figure 1, repetition neurons tend to become progressively more strongly activated as the repetition sequence continues.

We inspected repetition neurons in three English and one Japanese pre-trained language model. Our experimental results show that repetition neurons appear in both intermediate and final layers (§3.2). Furthermore, we demonstrate that deactivating these neurons suppresses the output probabilities of repeated tokens (§4.1), while activating them increases these probabilities (§4.2)¹. In addition, we highlight the relationship between repetition neurons and induction heads (§5).

¹Code for our experiments is available at https://github.com/tatHi/repetition_neuron

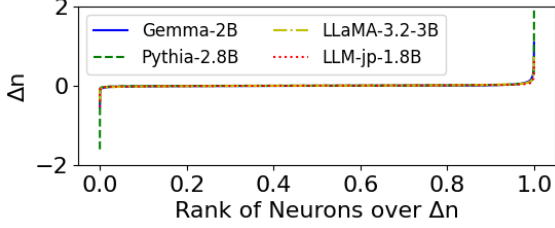


Figure 2: Δ_n of all neurons sorted in the ascending order. The x-axis shows the relative rank of each neuron (i.e., 1.0 is the 294,912-th neuron in Gemma-2B).

2 General Setting

2.1 Models

We utilized three English pre-trained language models: Gemma-2B (Team et al., 2024), which has 18 layers (294,912 neurons), Pythia-2.8B-Deduped (Biderman et al., 2023), with 32 layers (327,680 neurons), and LLaMA-3.2-3B (Dubey et al., 2024), with 28 layers (229,376 neurons). Additionally, we employed a Japanese pre-trained language model: LLM-jp-3-1.8B (LLM-jp et al., 2024), which has 24 layers (172,032 neurons).

2.2 Dataset with Repetition

To analyze the internal workings of language models on repetitive text, we collected 1,000 texts containing repetition from each language model. We randomly generated the first ten tokens with temperature = 1.0 using the generate() method from HuggingFace Transformers (Wolf et al., 2020). Afterward, we filtered out texts that did not contain repetition. We defined a text as containing repetition if the same 10-gram token sequence appeared three times at equal intervals within 100 tokens. Additionally, we excluded texts that did not have at least 50 tokens before and after the onset of repetition. The onset of repetition is defined as the point where the repeated sequence appears for the second time (see Figure 1). Table 2 in Appendix B provides examples of the repetitive texts generated through this process. The entire generation process took less than two hours on a single NVIDIA V100 GPU.

3 Finding Neurons Invoking Repetition

3.1 Detecting Repetition Neuron

In this work, we consider the outputs of the activation function in the feed-forward network of each Transformer layer as “neurons,” following previous studies (Geva et al., 2021; Dai et al., 2022; Wang

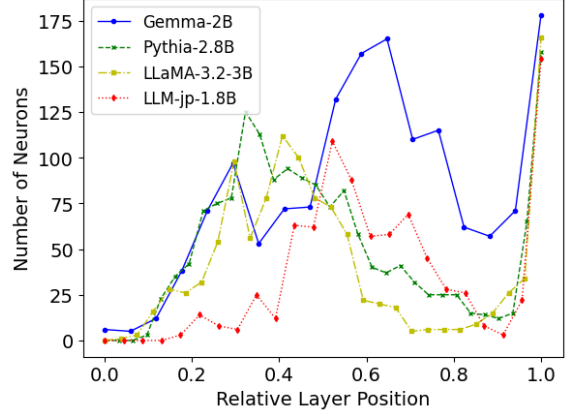


Figure 3: The number of repetition neurons for each layer when considering the top 0.5% of the entire neurons are repetition neurons. The x-axis shows the relative location of layers against the number of entire layers (e.g., 1.0 is the 18th layer in the case of Gemma-2B).

et al., 2022). We hypothesize that repetition neurons are more strongly activated in the range of texts with repetition and less active in texts without repetition. Therefore, we identify repetition neurons by comparing their activation values before and after the onset of repetition.

Let $x \in X$ represent a single text containing repetition, generated as described in §2.2, with $|X| = 1,000$ texts in total. Each text consists of a sequence of M tokens, $x = \{w_1, \dots, w_m, \dots, w_M\}$, and includes a repetition onset point s . This means the sequence after position s (i.e., $x_{s \leq m} = \{w_{m=s}, \dots, w_M\}$) consists of repeated tokens. We define the r tokens preceding the onset point, $x_{s-r}^{s-1} = \{w_{s-r}, \dots, w_{s-1}\}$, as the “normal” range without repetition, and the r tokens following the onset point, $x_s^{s+r-1} = \{w_s, \dots, w_{s+r-1}\}$, as the “repetition” range. We used the hyperparameter $r = 30$ for the main experiments, and Appendix E reports the ablation study. For each neuron n involved in the forward computation of the language model, we compute the average activation values a_n and \bar{a}_n over both the normal and repetition ranges, respectively.

$$a_n = \frac{1}{|X| \times r} \sum_{x \in X} \sum_{m=s-r}^{s-1} f(w_m, x_1^m, n), \quad (1)$$

$$\bar{a}_n = \frac{1}{|X| \times r} \sum_{x \in X} \sum_{m=s}^{s+r-1} f(w_m, x_1^m, n), \quad (2)$$

where $f(w_m, x_1^m, n)$ is a function that returns the activation value of neuron n at the time step corresponding to the input token w_m when reading the

sequence x_1^m with the language model. Next, we calculate the difference Δ_n between the activation values in the normal and repetition ranges as a score to quantify the effect of neurons on repetition:

$$\Delta_n = \bar{a}_n - a_n. \quad (3)$$

Here, larger Δ_n means the neuron n are activated more strongly in the repetition range than the normal range. We define the top K neurons with the largest Δ_n as repetition neurons for the model θ .

3.2 Observation of Repetition Neuron

Figure 2 shows the obtained Δ_n of all neurons, sorted in ascending order, for four language models. It is evident that only a small number of neurons exhibit remarkably high Δ_n values. This distribution is consistent with existing reports, which suggest that neuron activation is typically sparse (Li et al., 2023b; Voita et al., 2023). This also indicates that only a limited number of repetition neurons are activated exclusively in the repetition range.

Figure 1 shows the average activation values of the top four repetition neurons in Gemma-2B, measured across 1,000 texts for 30 tokens before and after the beginning of repetition. As repetition continues, the activation values of these neurons increase. This finding suggests that the repetition neurons respond to the recurrence of input tokens. We hypothesize that when the repetition neurons are strongly activated, the model starts to interpret copying previous tokens as a task, thereby falling into repetition (see §4).

Figure 3 presents the distribution of repetition neurons across different layers. The last layer contains the largest number of repetition neurons in all models, while a secondary peak appears in the intermediate layers. This suggests the existence of two types of repetition neurons: those that detect repeating patterns in the intermediate layers and those that drive the model to replicate previous contexts in the uppermost layer. The presence of repetition neurons in both the final and intermediate layers aligns with previous findings that task-specific neurons tend to reside in higher layers (Wang et al., 2022), and task-related parameters and hidden states are often found in intermediate layers (Hendel et al., 2023; Merullo et al., 2024). Figure 9a and 9b in Appendix E show that the location patterns of the repetition neurons remain consistent across variations in hyperparameters $|X|$ and r .

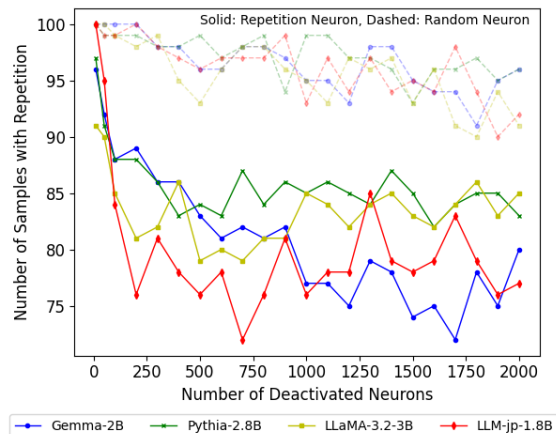


Figure 4: The number of samples with repetition after deactivating the repetition neurons for the texts originally with repetition.

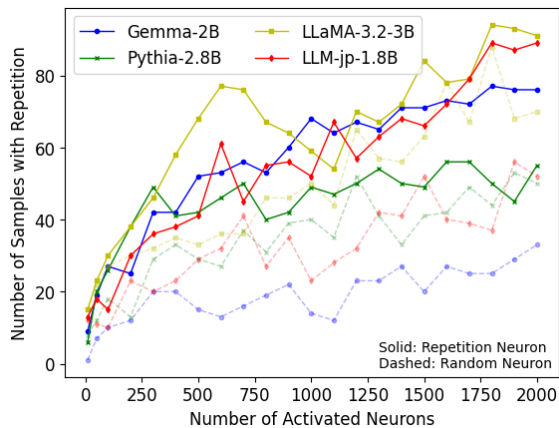


Figure 5: The number of samples with repetition after activating the repetition neurons for the texts originally without repetition.

4 Intervention to Repetition Neurons

If our hypothesis that the repetition neurons invoke repetition is correct, we should be able to control the repetition problem by intervening with these neurons (Arora et al., 2018; Wang et al., 2022).

4.1 Preventing Repetition

Setup: In this section, we test whether deactivating the repetition neurons can more effectively suppress repetition compared to deactivating randomly selected neurons. For this experiment, we generated an additional unseen 100 texts containing repetition for each language model, using the same method described in §2.2. We then deactivated the repetition neurons by setting their activation values to 0.0, starting from the token where the original text begins to repeat (e.g., in the case of Figure 1, from the generation step of second ‘‘H’’).

Original Greedy Output	Intervened Greedy Output
<p>The latest trend in design for the kitchen sink drain is the use of a stainless steel sink drain. This is a great way to add a touch of class to your kitchen. Stainless steel sinks are also very durable and easy to clean. The ▷ stainless steel sink drain is a great way to add a touch of class to your kitchen. It is also very durable and easy to clean. What is a stainless steel sink drain? A stainless steel sink drain is a type of sink drain that is made from stainless steel. Stainless steel is a type of metal that is resistant to corrosion and rust.</p>	<p>The latest trend in design for the kitchen sink drain is the use of a stainless steel sink drain. This is a great way to add a touch of class to your kitchen. Stainless steel sinks are also very durable and easy to clean. The ► stainless steel sink drain is a great way to add a touch of class to your kitchen. Stainless steel sinks are also very durable and easy to clean. The stainless steel sink drain is a great way to add a touch of class to your kitchen. Stainless steel sinks are also very durable and easy to clean. The stainless steel sink drain ...</p>

Table 1: The example of generation by Gemma-2B with and without intervention to the repetition neurons. ► indicates the beginning point of the intervention to invoke the repetition. We also indicate this point in the original greedy output with ▷ for visibility. Color-boxes show the repeating phrases.

Result: Figure 4 shows the number of samples containing repetition after deactivating varying numbers of repetition neurons (solid lines) compared to randomly selected neurons (dashed lines). As the figure demonstrates, deactivating the repetition neurons effectively reduces the number of samples with repetition compared to deactivating randomly selected neurons. This result confirms that the repetition neurons identified by our method are indeed responsible for causing the repetition problem. We observed that deactivating repetition neurons reduces the number of samples with repetition by up to 25% (and by as much as 35% with optimal hyperparameter settings, as shown in Figures 10a and 10b). This suggests that roughly 30% of the repetition problem can be attributed to the repetition neurons. Table 3 in §C provides an example where repetition was successfully suppressed, illustrating that the generated text remains grammatically coherent despite neuron intervention. Besides, the perplexity is not largely damaged by deactivating the repetition neurons, as shown in Figure 8a, which supports the coherency of the performance quantitatively. This confirms that the repetition neurons are specifically responsible for triggering repetition.

4.2 Invoking Repetition

Setup: In contrast to the experiment in §4.1, this section investigates whether activating the repetition neurons leads the model to produce repetitive outputs more effectively than activating randomly selected neurons. We newly prepared 100 unseen samples for each language model that do not contain repetition. Each sample consists of 210 tokens, with the first 10 tokens generated randomly and the remaining tokens generated greedily. Similar to the experiments in §4.1, we forcibly activate the repe-

tion neurons starting from the 51st token during the generation process. The neurons are activated by adding 1.0 to their original activation values.

Result: Figure 5 presents the number of samples exhibiting repetition after activating repetition neurons and randomly selected neurons. The figure demonstrates that repetitive samples increase as more neurons are activated. Furthermore, the activation of repetition neurons is more effective at invoking repetition compared to the activation of randomly selected neurons. Figure 8b also demonstrates that activating repetition neurons significantly worsens perplexity, suggesting an increased likelihood of generating repetitive tokens. These results support our hypothesis that neurons with higher Δ_n function as “skill neurons” that trigger repetitive behavior. Activating randomly selected neurons also leads to many repetitive samples, suggesting that factors like unstable hidden states also contribute to the repetition problem in addition to the repetition neurons.

Case Study: Table 1 provides a typical generation example obtained by activating the repetition neurons. The table highlights the text range where we forcibly activate the repetition neurons with the bold font. Interestingly, the model does not immediately begin repeating tokens following the intervention. Instead, once it completes the sentence it is generating, the model starts to replicate text that appeared **before** the point of intervention. This suggests that the repetition neurons encourage the model to copy previous outputs rather than simply generating tokens that are easily repeated.

5 Comparison with Induction Heads

Several works in in-context learning have examined how attention heads, particularly induction heads (Olsson et al., 2022; Bansal et al., 2023;

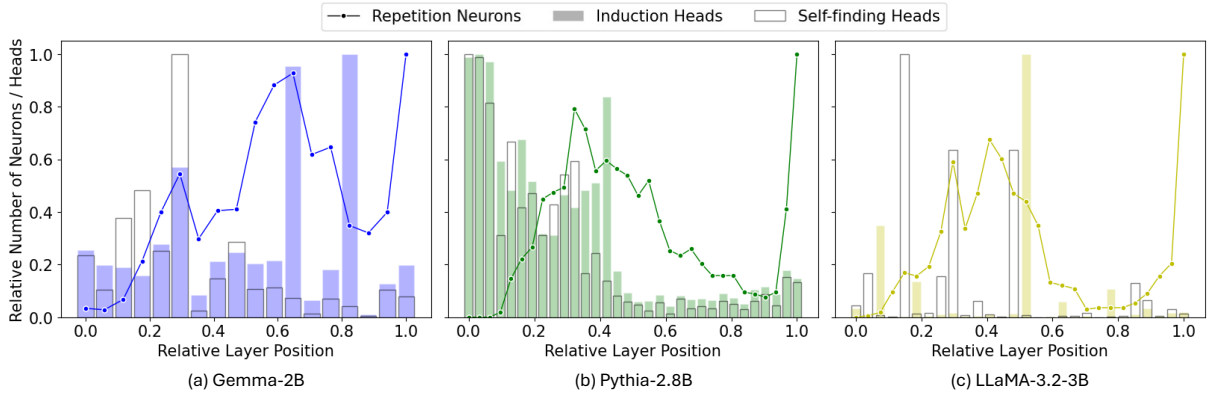


Figure 6: Frequency of reputation neurons (lines), induction heads (colored bars), and self-finding heads (edged bars) for repetition over three English models.

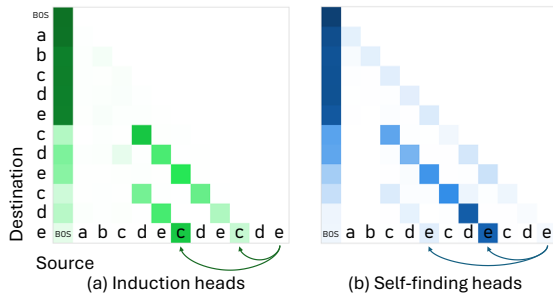


Figure 7: Examples of attention heat-maps for induction and self-finding heads capturing repetition inputs: “a b c d e c d e c d e”. The expected next token is “c”.

Crosbie and Shutova, 2024), exhibit repetitive behaviors. These studies explore how in-context learning performance relates to an LLM’s ability to copy patterns in synthetically generated repeating sequences. Building on these insights, we focus on a neuron-based analysis of repetition actually generated by the LLMs themselves. In this section, we compare the “repetition neurons,” “induction heads,” and “self-finding heads” derived from the same repeating texts using three LLMs: Gemma-2B, Pythia-2.8B, and LLaMA-3.2-3B.

Figure 6 shows the distribution of the repetition neurons (the same results as in Figure 3), induction heads, and self-finding heads over layers. In our analysis, we define “induction heads” as heads attending to repeating tokens that are to be generated after the current input token (Figure 7a). We also define “self-finding heads” as heads attending to the repeating token identical to the input token (Figure 7b). We identified a head as induction or self-finding if its total attention score for the target tokens that appear after the second repeating position exceeds 0.5. We then summarize their

layer positions to see how these heads align with repetition neurons.

The observed behavior varies by model architecture. As shown in Figure 6a, Gemma-2B’s repetition neurons share two peaks with the induction heads, and one of these peaks is also shared by self-finding heads. This suggests that certain repetition neurons are activated in response to both induction and self-finding heads capturing repetition. However, the highest induction-head peak (layer 14 of 18) does not coincide with the highest repetition-neuron peak (layer 18 of 18).

Figures 6b and 6c present a different pattern for Pythia and LLaMA, where we do not observe a strong alignment between repetition neurons and induction heads. Nevertheless, similar to Gemma-2B, some peaks in the early layers of repetition neurons correspond to peaks of self-finding heads. This suggests that repetition neurons respond to self-finding patterns in earlier layers and take on different roles in later layers.

Overall, this comparison among repetition neurons, induction heads, and self-finding heads reveals coordinated interactions while showing their distinct roles in detecting and invoking repetition.

6 Conclusion

We proposed a method to identify the repetition neurons that contribute to the repetition problem in text generation. These neurons are located in both the intermediate and final layers of the Transformer, similar to skill neurons and task vectors. Our experimental results show that by intervening in the activity of these repetition neurons, we can control the occurrence of repetitive outputs.

Limitations

The primary goal of this short paper is to report the existence of repetition neurons in repetitive texts and to describe their basic behavior. We recognize that our findings are likely to spark further discussion, which lies beyond the scope of this work. To facilitate future research, we outline several key topics related to repetition neurons:

- We prepared the dataset without considering detailed aspects of repetition (§2.2), such as the length of each repetitive phrase. By focusing on specific phrase lengths, can we identify particular tendencies in the behavior of repetition neurons?
- We observed two distinct peaks in the distribution of repetition neurons across layers in Figure 3. What are the functional differences between neurons located in the intermediate layers and those in the final layer?
- The experimental results of deactivating the repetition neuron suggest that roughly 30% of the repetition problems are caused by the repetition neuron (§4.1). What causes the rest 70% repetition problem?
- Does the behavior of repetition neurons change against the model configuration (e.g., the parameter size, the language used in the pre-training, the activation functions, and so on)?
- We used the simple intervention to the repetition neurons: replacing the activation value with 0.0 for deactivation (§4.1) and adding 1.0 for activation (§4.2). What can we observe when gradually increasing or decreasing the activation value instead of the simple replacement or addition?
- Beyond the neuron-based and head-based analysis (§5), can we find any other specific circuit in the LLMs' calculation when outputting the repetitive texts?

Some of the above topics are partially discussed in the appendix. We believe that our findings in this paper help the further discussion to reveal the inner working of the repetition problem.

Acknowledgement

This work was supported by JST, CREST Grant Number JPMJCR20D2, Japan.

References

- Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. 2018. Stronger generalization bounds for deep nets via a compression approach. In *International conference on machine learning*, pages 254–263. PMLR.
- Hritik Bansal, Karthik Gopalakrishnan, Saket Dingliwal, Sravan Bodapati, Katrin Kirchhoff, and Dan Roth. 2023. Rethinking the role of scale for in-context learning: An interpretability-based case study at 66 billion scale. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. Preprint, arXiv:2005.14165.
- Yuheng Chen, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. 2024. Journey to the center of the knowledge neurons: Discoveries of language-independent knowledge neurons and degenerate knowledge neurons. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17817–17825.
- Joy Crosbie and Ekaterina Shutova. 2024. Induction heads as an essential mechanism for pattern matching in in-context learning. *arXiv preprint arXiv:2407.07011*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman,

- Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. [arXiv preprint arXiv:2407.21783](#).
- Zihao Fu, Wai Lam, Anthony Man-Cho So, and Bei Shi. 2021. A theoretical analysis of the repetition problem in text generation. In [Proceedings of the AAAI Conference on Artificial Intelligence](#), volume 35, pages 12848–12856.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In [Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing](#), pages 5484–5495.
- Roei Hendel, Mor Geva, and Amir Globerson. 2023. In-context learning creates task vectors. In [Findings of the Association for Computational Linguistics: EMNLP 2023](#), pages 9318–9333.
- Maor Ivgi, Ori Yoran, Jonathan Berant, and Mor Geva. 2024. From loops to oops: Fallback behaviors of language models under uncertainty. [arXiv preprint arXiv:2407.06071](#).
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. [arXiv preprint arXiv:1909.05858](#).
- Huayang Li, Tian Lan, Zihao Fu, Deng Cai, Lemao Liu, Nigel Collier, Taro Watanabe, and Yixuan Su. 2023a. Repetition in repetition out: Towards understanding neural text degeneration from the data perspective. [Advances in Neural Information Processing Systems](#), 36:72888–72903.
- Zonglin Li, Chong You, Srinadh Bhojanapalli, Daliang Li, Ankit Singh Rawat, Sashank J. Reddi, Ke Ye, Felix Chern, Felix Yu, Ruiqi Guo, and Sanjiv Kumar. 2023b. [The lazy neuron phenomenon: On emergence of activation sparsity in transformers](#). In [The Eleventh International Conference on Learning Representations](#).
- LLM-jp, :, Akiko Aizawa, Eiji Aramaki, Bowen Chen, Fei Cheng, Hiroyuki Deguchi, Rintaro Enomoto, Kazuki Fujii, Kensuke Fukumoto, Takuya Fukushima, Namgi Han, Yuto Harada, Chikara Hashimoto, Tatsuya Hiraoka, Shohei Hisada, Sosuke Hosokawa, Lu Jie, Keisuke Kamata, Teruhito Kanazawa, Hiroki Kanezashi, Hiroshi Kataoka, Satoru Katsumata, Daisuke Kawahara, Seiya Kawano, Atsushi Keyaki, Keisuke Kiryu, Hirokazu Kiyomaru, Takashi Kodama, Takahiro Kubo, Yohei Kuga, Ryoma Kumon, Shuhei Kurita, Sadao Kurohashi, Conglong Li, Taiki Maekawa, Hiroshi Matsuda, Yusuke Miyao, Kentaro Mizuki, Sakae Mizuki, Yugo Murawaki, Ryo Nakamura, Taishi Nakamura, Kouta Nakayama, Tomoka Nakazato, Takuro Niitsuma, Jiro Nishitoba, Yusuke Oda, Hayato Ogawa, Takumi Okamoto, Naoaki Okazaki, Yohei Oseki, Shintaro Ozaki, Koki Ryu, Rafal Rzepka, Keisuke Sakaguchi, Shota Sasaki, Satoshi Sekine, Kohei Suda, Saku Sugawara, Issa Sugiura, Hiroaki Sugiyama, Hisami Suzuki, Jun Suzuki, Toyotaro Suzumura, Kensuke Tachibana, Yu Takagi, Kyosuke Takami, Koichi Takeda, Masashi Takeshita, Masahiro Tanaka, Kenjiro Taura, Arseny Tolmachev, Nobuhiro Ueda, Zhen Wan, Shuntaro Yada, Sakiko Yahata, Yuya Yamamoto, Yusuke Yamauchi, Hitomi Yanaka, Rio Yokota, and Koichiro Yoshino. 2024. [LLM-jp: A cross-organizational project for the research and development of fully open japanese llms](#).
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. 2024. [Circuit component reuse across tasks in transformer language models](#). In [The Twelfth International Conference on Learning Representations](#).
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. [arXiv preprint arXiv:2209.11895](#).
- Alec Radford, Rafal Jozefowicz, and Ilya Sutskever. 2017. Learning to generate reviews and discovering sentiment. [arXiv preprint arXiv:1704.01444](#).
- 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.
- Keisuke Shirai, Kazuma Hashimoto, Akiko Eriguchi, Takashi Ninomiya, and Shinsuke Mori. 2021. Neural text generation with artificial negative examples to address repeating and dropping errors. [Journal of Natural Language Processing](#), 28(3):751–777.
- Yifan Song, Guoyin Wang, Sujian Li, and Bill Yuchen Lin. 2024. The good, the bad, and the greedy: Evaluation of llms should not ignore non-determinism. [arXiv preprint arXiv:2407.10457](#).
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepey, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni,

- Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. [Gemma: Open models based on gemini research and technology](#).
- Aditya Vaidya, Javier Turek, and Alexander Huth. 2023. Humans and language models diverge when predicting repeating text. In [Proceedings of the 27th Conference on Computational Natural Language Learning \(CoNLL\)](#), pages 58–69.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In [Advances in Neural Information Processing Systems](#), volume 30. Curran Associates, Inc.
- Elena Voita, Javier Ferrando, and Christoforos Nalmpantis. 2023. Neurons in large language models: Dead, n-gram, positional. [arXiv preprint arXiv:2309.04827](#).
- Weichuan Wang, Zhaoyi Li, Defu Lian, Chen Ma, Linqi Song, and Ying Wei. 2024. Mitigating the language mismatch and repetition issues in llm-based machine translation via model editing. [arXiv preprint arXiv:2410.07054](#).
- Xiaozhi Wang, Kaiyue Wen, Zhengyan Zhang, Lei Hou, Zhiyuan Liu, and Juanzi Li. 2022. Finding skill neurons in pre-trained transformer-based language models. In [Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing](#), pages 11132–11152.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In [Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations](#), pages 38–45, Online. Association for Computational Linguistics.
- Jin Xu, Xiaojiang Liu, Jianhao Yan, Deng Cai, Huayang Li, and Jian Li. 2022. Learning to break the loop: Analyzing and mitigating repetitions for neural text generation. [Advances in Neural Information Processing Systems](#), 35:3082–3095.
- Jianhao Yan, Jin Xu, Chiyu Song, Chenming Wu, Yafu Li, and Yue Zhang. 2024. [Understanding in-context learning from repetitions](#). In [The Twelfth International Conference on Learning Representations](#).
- Wenhong Zhu, Hongkun Hao, and Rui Wang. 2023. Penalty decoding: Well suppress the self-reinforcement effect in open-ended text generation. In [Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing](#), pages 1218–1228.

A Comparison with Existing Work

The concept of the “repetition neuron” was first introduced by Wang et al. (2024), where they showed its impact on machine translation using in-context learning. While our research was conducted concurrently and independently, the fact that multiple research teams are exploring similar topics underscores the growing interest in understanding repetition within NLP models. Our findings in general text generation complement their results in the machine translation task, supporting the idea that repetition neurons play a broader role across various generation tasks. Below, we outline the key distinctions between our work and theirs to highlight our unique contributions.

Unlike their focus on improving performance in in-context learning for machine translation by editing repetition neurons, our research aims to uncover the inner workings of LLMs when processing repetitive text, specifically from the perspective of repetition neurons. To achieve this, we employed four different pre-trained language models (three English, one Japanese) and demonstrated that repetition neurons are not restricted to a single architecture on a specific task like machine translation with LLaMA-7B but are observable across various architectures (§2.1, 2.2). Our broader focus on general text generation highlights the versatility of the repetition neuron phenomenon, as compared to the task-specific nature of the machine translation context used in Wang et al. (2024).

Our experiments also provide a more detailed analysis of the distribution of repetition neurons across layers in different models and under varying hyperparameters, something not covered in previous work (§3.2 and §E). While both studies involve deactivating repetition neurons to observe the impact on generation, our experiments present a comprehensive comparison across four language models, revealing performance changes as the number of deactivated neurons varies (§4.1). One insight that emerges from our findings is that selecting only the top 300 neurons, as in Wang et al. (2024), may be insufficient for models of larger scale, a point we explore in depth. In addition, our exploration of neuron activation to deliberately induce repetition (§4.2) introduces a novel dimension to this research.

Methodologically, our approach to identifying repetition neurons by comparing activation values before and after the repetition point is more

straightforward than their attribution score-based method (Dai et al., 2022). Given that both methods yield similar outcomes in terms of controlling repetition, our simpler approach could serve as an alternative for identifying repetition neurons in large-scale models.

In sum, while our findings do not conflict with those of Wang et al. (2024), our work complements their research by providing a broader, more detailed exploration of the role of repetition neurons. Our findings not only validate the existence of these neurons across different architectures but also contribute novel insights into their layer-wise distribution and activation patterns. These insights pave the way for more targeted interventions in controlling repetition across various language generation tasks.

B Example of Generated Repetition

Table 2 shows the actual examples of the dataset created in the manner explained in §2.2. In this table, the bold font highlights the repeated phrases. Note that here we highlight the text range from the first repeating phrases to the end of the third repeating phrase, while the repetition range mentioned in §3.1 refers to the span after the second repeating point. As shown in this table, there are various lengths of repeated phrases in each sample. Future work should focus on the effect of differences in the repetition style on the repetition neurons.

C Additional Case Study

Table 3 shows an example of text generation with Gemma-2B, where we deactivate the top 600 repetition neurons. The original greedy generation falls into the repetition of the short phrase “The driveway?” after listing similar phrases. By deactivating the repetition neurons, the model terminates to list similar phrases and begins to generate natural sentences. This result implies that some repetition neurons have an effect of making the model copy the template “the ___?” and some other neurons are in charge of copying “driveway”.

D Perplexity with Intervention

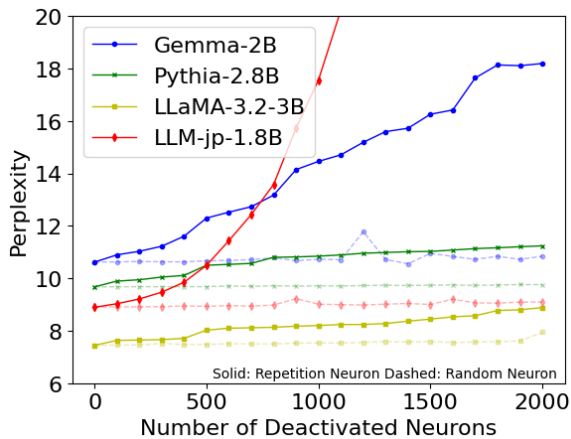
Figure 8a and 8b show the changes in perplexity based on different numbers of intervened repetition neurons. We used the test split of WikiText-2 for all models, including LLM-jp-1.8B, which was fine-tuned on Japanese corpora. As shown in Figure

<p><i>Gemma-2B</i> ... ask this, but I'm curious about the black hat. I've been looking for a good one for a while now, but I've never found one that I really liked. I've been looking at the black hat, but I'm not sure if it's the right one for me. I'm not sure if it's the right one for me. I'm not sure if it's the right one for me. I'm not sure if it's the right one for me. come to me?" Smith said she has seen a rise in the number of students who are dealing with mental health issues. "I think it's because of the pandemic," she said. "I think it's because of the isolation. I think it's because of the stress of the world. I think it's because of the stress of the world. I think it's because of the stress of the world I have been using the same starter for the last 2 months. I have been using the same flour and water. I have been using the same method. I have been using the same oven. I have been using the same technique. I have been using the same everything. I have been using the same everything. I have been using the same everything. I ...</p>
<p><i>Pythia-2.8B-Deduped</i> ... good. I did add a little more salt than the recipe called for. I also added a little more pepper. I also added a little more garlic powder. I also added a little more onion powder. I also added a little more oregano. I also added a little more basil. I also added a little more basil. I also added a little more basil a bad mood. "I'm sorry," I said. "I didn't mean to upset you." "It's okay," she said. "I'm just glad you're here." "I'm glad I'm here, too." "I'm glad you're here, too." "I'm glad you're here, too." "I'm glad you're here, too." "I'm glad you're here, too." look at the flower and see if it is a little bit different. If it is, you should not be too concerned. If it is a little bit different, you should be able to use it. If you are using a new flower, you should be able to use it. If you are not, you should be able to use it. If you are not, you should be able to use it. If you are not ...</p>
<p><i>LLaMA-3.2-3B</i> ... It is a book of the plan of salvation. It is a book of the plan of happiness. It is a book of the plan of mercy. It is a book of the plan of redemption. It is a book of the plan of life. It is a book of the plan of happiness. It is a book of the plan of happiness. It is a book of the plan of happiness hesion and proliferation of human breast cancer cells. The CD44 expression was analyzed by flow cytometry and immunohistochemistry. The CD44 expression was also analyzed by reverse transcription-polymerase chain reaction (RT-PCR) and Western blotting. The CD44 expression was also analyzed by immunohistochemistry. The CD44 expression was also analyzed by immunohistochemistry. The CD44 expression was also analyzed by immunohistochemistry ? A) 10 B) 12 C) 15 D) 20 E) 25 Answer: Let x be the number of buses on the route. $21x/4 = 21x/4 * 3/4 = 3/4 * 21x = 3/4 * 21x = 3/4 * 21x ...$</p>
<p><i>LLM-<i>jp</i>-3-1.8B</i> ...と言つて豆まきをします。この豆まきは、「鬼」を追い払うという意味があります。「鬼」は「外」に追い出して、「福」は「内」に招き入れる。つまり、「鬼」を追い出して、「福」を招き入れる。「鬼」を追い出して、「福」を招き入れる。「鬼」を追い出して、「福」を招き入れる。... ...第5回: 第4回の続き ## 第4回: 第3回の続き ## 第3回: 第2回の続き ## 第2回: 第1回の続き ## 第1回: 第0回の続き ## 第0回: 第0回の続き ## 第0回: 第0回の続き ## 第0、必要な措置を講ずることを命ずることができる。(貸貸人の責めに帰すべき事由による場合の契約解除) 第十三条 貸貸人は、貸貸借の期間が満了する前に貸貸借の期間が満了する前に ...</p>

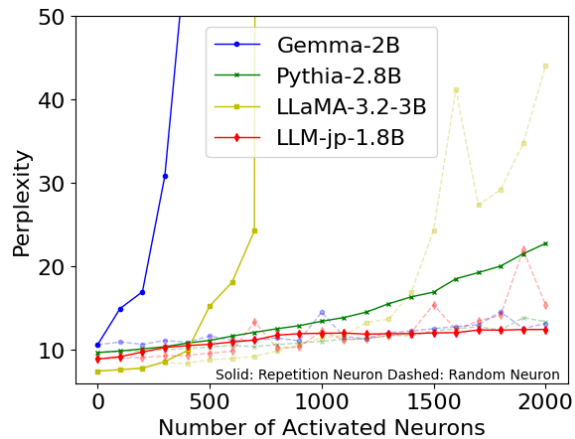
Table 2: Three examples of generated samples with perturbation by each language model. This table shows 30 tokens before the beginning of repetition and three times repetition.

Original Greedy Output	Intervened Greedy Output
What's the one thing you love the most about your home? Is it the view? The location? The size? The layout? The design? The style? The architecture? The materials? The colors? The finishes? The fixtures? The appliances? The lighting? The furniture? The art? The decor? The landscaping? The pool? The fireplace? The kitchen? The bathroom? The closet? The storage? The garage? The driveway? ...	What's the one thing you love the most about your home? Is it the view? The location? The size? The layout? The design? The style? The architecture? The materials? The colors? The finishes? The fixtures? The appliances? The lighting? The furniture? The art? The decor? The landscaping? The pool? The fireplace? The kitchen? The bathroom? The closet? The storage? The garage? The driveway? If you're like most of us, you probably have a million and one things you love about your home. But, if you're like me, you also have a million and one things you don't like about your home. I've lived in my 1950s-era home for ...

Table 3: An example of text generation that repetition is reduced by deactivating repetition neurons (Gemma-2B). The repeated texts are highlighted with underline in the original greedy output. In the intervened output, bold indicates the range where the repetition neurons are deactivated.



(a) Deactivation



(b) Activation

Figure 8: Perplexity on the test split of WikiText-2 with the intervention to repetition neurons.

8a, the performance degradation caused by deactivating repetition neurons is relatively moderate. Although the impact on perplexity is smaller when intervening on randomly sampled neurons, considering that GPT-2’s perplexity on the same corpus was 18.34 (Radford et al., 2019), the degradation caused by deactivating repetition neurons is acceptable. This suggests that repetition neurons do not significantly affect the generation of normal texts that do not contain repetition.

Unlike the English models, LLM-jp-1.8B’s perplexity increases substantially even when a smaller number of repetition neurons are deactivated. This result implies that repetition neurons may be language-specific. In other words, neurons identified as repetition neurons in Japanese may serve a different role in English texts, leading to more significant harm to perplexity on English test sets.

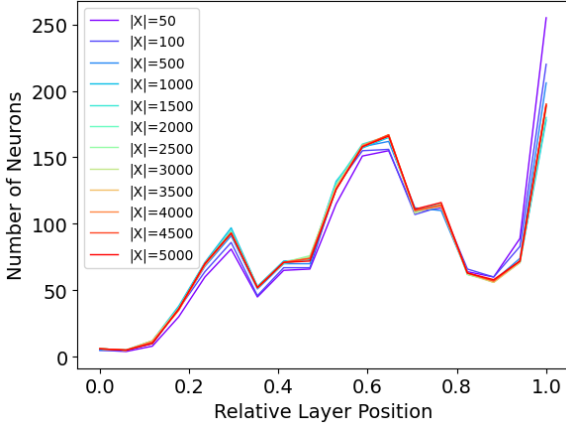
In contrast, the perplexity increases dramatically when repetition neurons are activated. For instance, the perplexity of Gemma-2B and LLaMA-3.2-3B exceeds 100 when 500 and 800 repetition neurons are activated, respectively, indicating that the model becomes severely impaired with the activation of a large number of these neurons. This suggests that repetition neurons play an important role in generating non-grammatical outputs, and their improper activation increases the likelihood of tokens reappearing from earlier in the text. On the other hand, LLM-jp-1.8B’s perplexity remains largely unaffected by activating its repetition neurons. This further suggests that repetition neurons could be language-specific, as those found in Japanese texts do not have a significant impact on the perplexity of English texts.

E Ablation Study

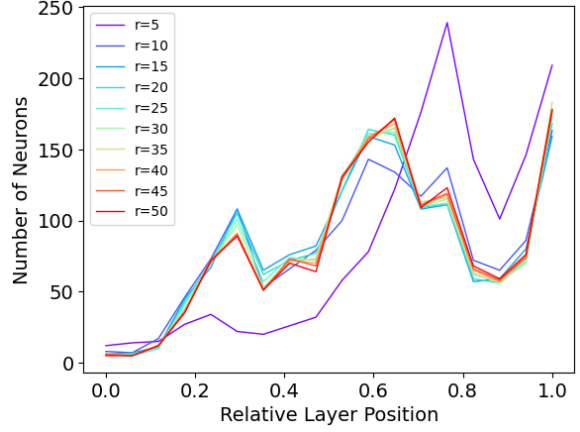
The proposed method to seek the repetition neurons has two hyperparameters: the number of repetitive texts generated for the dataset $|X|$ (§2.2) and the text range r to be focused on when calculating activation scores (§3.1). In the main body of this paper, we used $|X| = 1,000$ and $r = 30$. Herein, we investigate the effect of these hyperparameters on the same experiments using Gemma-2B. The scope of this ablation study is $|X| = \{50, 100, 500, 1000, 1500, \dots, 5000\}$ and $r = \{5, 10, 15, \dots, 50\}$. When investigating the various $|X|$, we fix the other hyperparameter as $r = 30$, while $|X| = 1,000$ for the investigation of r .

Figure 9a and 9b show the location of the repetition neuron on the Transformer layers (§3.2). As shown in the figure, the size of the dataset $|X|$ does not have large effect on the distribution, which means we can obtain the similar set of repetition neurons both with smaller and larger sizes of datasets. On the other hand, r has an effect on the distribution to some degree. For the case of Gemma-2B, we can obtain roughly a similar tendency with $15 \leq r$.

Figure 10a and 10b show the difference in the performance for the experiment about deactivating repetition neurons (§4.1). The experimental result indicates that we can obtain the more reduction effect with the larger number of deactivated repetition neurons. In contrast, the larger r cause the decrease of the effect to reduce the repetition with larger number of deactivated neurons. Figure 10b suggests that $r = 10$ or $r = 15$ leads to the largest

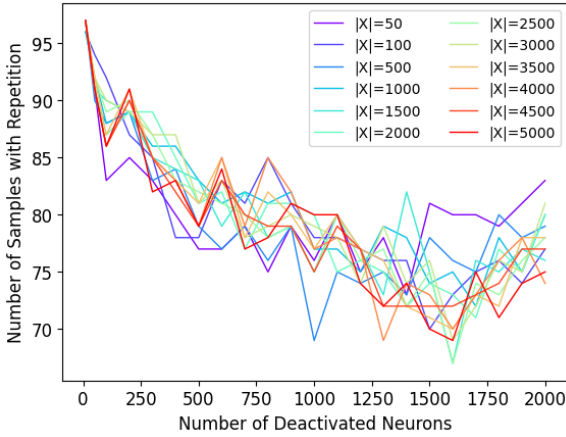


(a) Difference against $|X|$

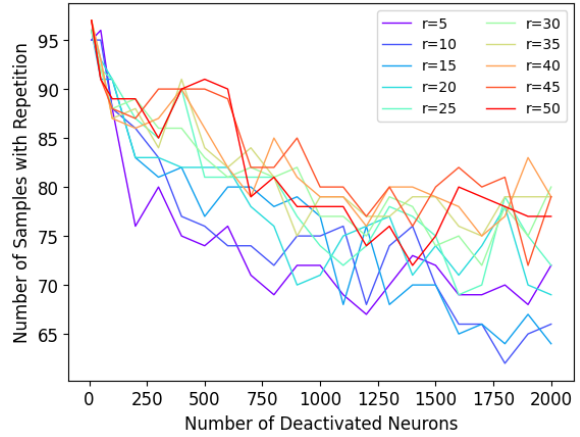


(b) Difference against r

Figure 9: The number of repetition neurons for each layer with various hyperparameters (Gemma-2B).



(a) Difference against $|X|$



(b) Difference against r

Figure 10: The experimental result with deactivating repetition neurons for different hyperparameters (Gemma-2B).

effect on controlling the repetitive generation.

The experimental results for the effect of various hyperparameters on the setting with activating repetition neurons in Figure 11a and 11b show similar trends. The larger $|X|$ leads to the larger number of repetitive texts while $r = 10$ or $r = 15$ has the largest effect on the repetitive generation. These results could be an important clue to investigate the relation between the length of repetitive phrases and the repetition neurons in the future work.

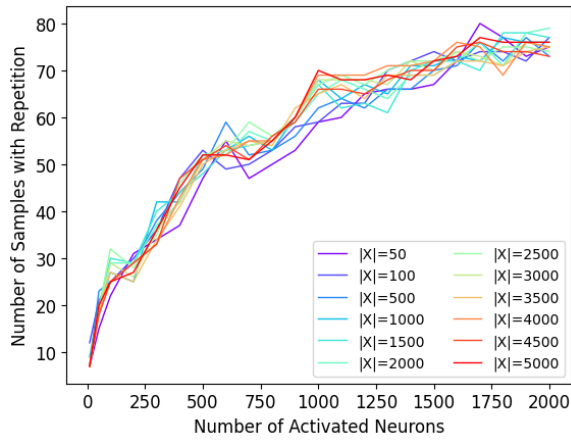
F Comparison between Two Model Sizes

We compared the experimental results with two different sizes of the same architecture: Gemma-2B and Gemma-7B.

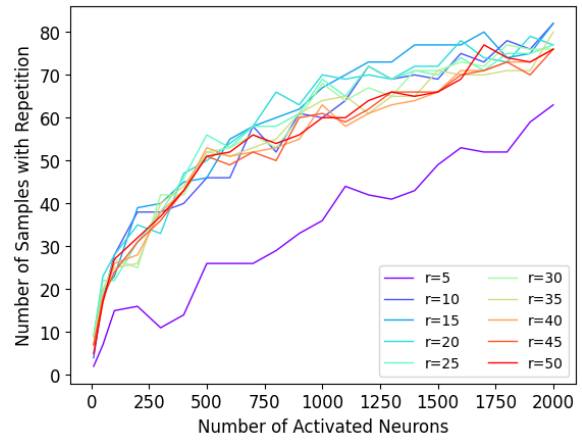
Figure 12 shows the distribution of repetition neurons over the layers. Compared to Gemma-2B, the repetition neurons of Gemma-7B are mainly located in the last layer. This result implies that

the nature of repetition neurons varies depending on the size of the language model instead of the architecture.

Figure 13 and 14 show the experimental results of the two experiments with deactivating and activating the repetition neurons, respectively. The results of Figure 13 indicate that the effect of repetition neurons to prevent the repetition problem becomes smaller when using the larger language model. This result aligns with the experiment shown in the existing work (Wang et al., 2024). In contrast, the activation of repetition neurons of Gemma-7B largely affects the repetitive outputs (Figure 14). These differences in performance show that there is room to be explored about the repetition neuron from broader viewpoints.



(a) Difference against $|X|$



(b) Difference against r

Figure 11: The experimental result with activating repetition neurons for different hyperparameters (Gemma-2B).

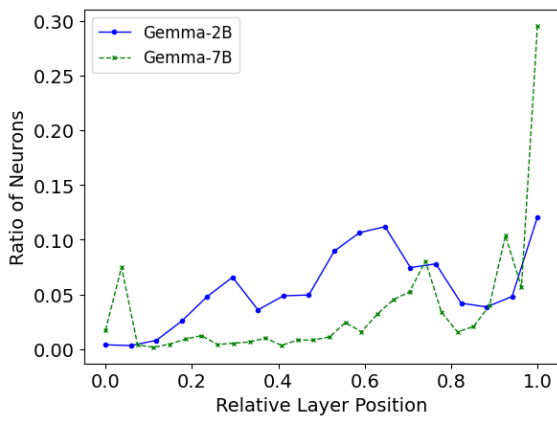


Figure 12: The number of top 0.5% repetition neurons for each layer. The x-axis shows the relative location of layers against the number of entire layers. The y-axis shows the relative number of neurons against the 0.5% neurons.

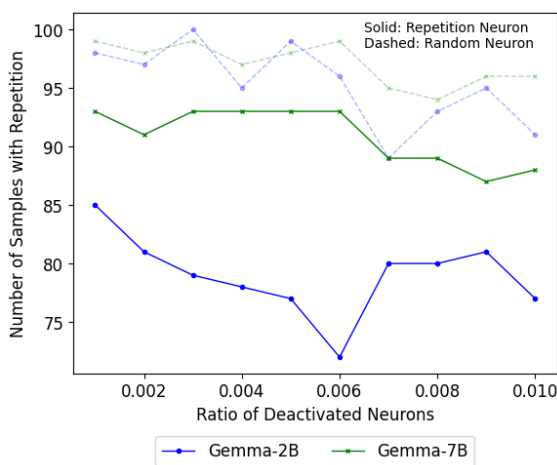


Figure 13: The experimental results with deactivating repetition neurons for two model sizes.

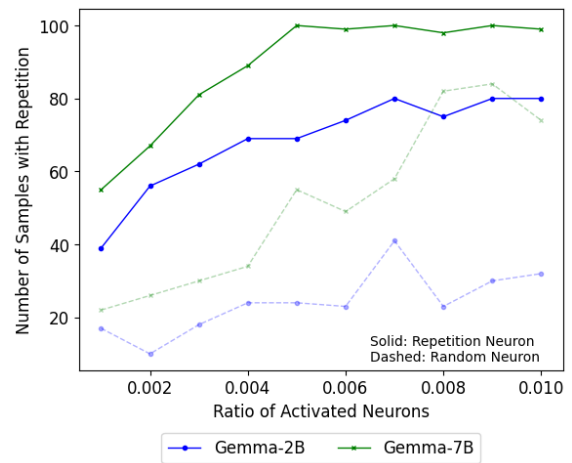


Figure 14: The experimental results with activating repetition neurons for two model sizes.