
Confidence Regulation Neurons in Language Models

Alessandro Stolfo*
ETH Zürich

Ben Wu*
University of Sheffield

Wes Gurnee
MIT

Yonatan Belinkov
Technion

Xingyi Song
University of Sheffield

Mrinmaya Sachan
ETH Zürich

Neel Nanda

Abstract

Despite their widespread use, the mechanisms by which large language models (LLMs) represent and regulate uncertainty in next-token predictions remain largely unexplored. This study investigates two critical components believed to influence this uncertainty: the recently discovered entropy neurons and a new set of components that we term token frequency neurons. Entropy neurons are characterized by an unusually high weight norm and influence the final layer normalization (LayerNorm) scale to effectively scale down the logits. Our work shows that entropy neurons operate by writing onto an *unembedding null space*, allowing them to impact the residual stream norm with minimal direct effect on the logits themselves. We observe the presence of entropy neurons across a range of models, up to 7 billion parameters. On the other hand, token frequency neurons, which we discover and describe here for the first time, boost or suppress each token’s logit proportionally to its log frequency, thereby shifting the output distribution towards or away from the unigram distribution. Finally, we present a detailed case study where entropy neurons actively manage confidence in the setting of induction, i.e. detecting and continuing repeated subsequences.

1 Introduction

As large language models (LLMs) increasingly permeate high-stakes applications, the lack of transparency in their decision-making processes poses significant vulnerabilities and risks [5]. Understanding the basis of these models’ decisions, especially how they regulate confidence in their predictions, is crucial not only for advancing model development but also for ensuring their safe deployment [1, 33]. LLMs have been empirically shown to be fairly well calibrated: on question-answering tasks, token-level probabilities of a model prediction generally match the probability of the model being correct [40, 58]. This raises the question of whether LLMs possess mechanisms for general-purpose calibration to mitigate the risks associated with overconfident predictions.

Significant research has been conducted on estimating model uncertainty in neural networks [19, 21] and specifically in LLMs [22]. While many studies focus on quantifying and calibrating model confidence [49, 40, 47, 45, 65, *inter alia*], there is little research into the internal mechanisms LLMs might use to calibrate their predictions.

In this work, we explore two types of components in Transformer-based language models that we believe serve a calibration function: the recently identified entropy neurons and a new class of components that we term *token frequency neurons*. Entropy neurons have been brought to light by recent work [42, 28] and are characterized by their high weight norm and their low composition with the unembedding matrix despite being in the final layer. The low composition with the unembedding matrix would suggest that they play a minor role in the next-token prediction.

*Equal contribution. Correspondence to stolfoa@ethz.ch, bpwu1@sheffield.ac.uk.

However, their high weight norm, despite LLMs being trained with weight decay [48], indicates that these neurons must be important for performance. This leads us to infer that these neurons may play a calibration role, as hypothesized by Gurnee et al. [28]. We show that entropy neurons work by writing to an effective null space of the unembedding matrix and leverage the layer normalization [4] that is applied to the residual stream. This way, these components effectively modulate the entropy of the model’s output distribution in a manner that only minimally affects the actual model prediction (Figure 1). We observe the presence of entropy neurons in GPT-2 [62], Pythia [6], Phi-2 [38], Gemma 2B [68], and LLaMA2 7B [70], demonstrating for the first time their role across different model families and scales.²

Motivated by the importance of the token frequency distribution (i.e., the unigram distribution) in language modeling [11, 51], we identify a novel class of neurons that boost or suppress each token’s logit proportionally to its frequency. We provide evidence showing that these neurons, which we term token frequency neurons, modulate the distance between the model’s output distribution and the unigram distribution, which the model defaults to in settings of high uncertainty.

As a case study, we analyze the activation of these neurons in scenarios involving induction (i.e., the repetition of subsequences in the input; 57). Our results show that entropy neurons increase the output distribution’s entropy, thus decreasing the model’s confidence in its predictions for repeated sequences. This represents a hedging mechanism, aimed at mitigating loss spikes on confidently wrong predictions.

Our contributions are as follows: (1) We show that models have dedicated circuitry to calibrate their confidence, with dedicated neurons in the final layer: the previously discovered entropy neurons and novel token frequency neurons. (2) We explore the mechanism by which each neuron family affects confidence (§3 and §4). Notably, we find that the model learns a low-rank effective null space for the unembedding that entropy neurons write to, and the final LayerNorm scale is used to calibrate the output logits. This striking phenomenon would have been missed by the assumptions made in standard methods such as direct logit attribution [16, 24, 46]. (3) We study how these neurons are used in practice to regulate confidence (§5). As expected, neurons that lower confidence worsen performance when models are correct and improve it when models are incorrect. We close with a mechanistic case study of how induction heads [57] use entropy neurons to control confidence when detecting and continuing repeated text (§6).³

2 Background

In this section, we provide an overview of the Transformer architecture [71], focusing on the components relevant to our analysis. Given a vocabulary $\mathcal{V}$, we denote an autoregressive language model as $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{Y}$ is the space of probability distributions over $\mathcal{V}$. $\mathcal{M}$ takes as input a token sequence $x = [x_1, \dots, x_m] \in \mathcal{X} \subseteq \mathcal{V}^m$, and outputs a probability distribution $P_{\text{model}} : \mathcal{V} \rightarrow [0, 1]$ to predict the next token in the sequence. In this work, we focus on decoder-only Transformer-based models, which represent the backbone of the most capable current AI systems [58, 26]. The Transformer architecture consists of two core components: the multi-head self-attention and a multi-layer perceptron (MLP). These two components read information in from and write out to the residual stream (i.e., the per-token hidden state consisting of the sum of all previous component outputs) [16].

Final-layer Neurons in LLaMA2 7B




Figure 1: **Entropy and Prediction.** We mean ablate final-layer neurons across 4000 tokens and measure the variation in the entropy of the model’s output P_{model} against average change of model’s prediction ($\text{argmax}_x P_{\text{model}}(x)$). We identify a set of neurons whose effect depends on LayerNorm (red points; metric described in §3.2), and which affect the model’s confidence (quantified as entropy of P_{model}) with minimal impact on the prediction.

²In the main paper, we report detailed results for GPT-2 Small and LLaMA2 7B. Results for other models are provided in the Appendix.

³Our code and data are available at <https://github.com/bpwu1/confidence-regulation-neurons>.

MLP. Central to this study is the structure of the MLP layers of the transformer.⁴ Given a normalized residual stream hidden state $\mathbf{x} \in \mathbb{R}^{d_{\text{model}}}$, the output of an MLP layer is

$$\text{MLP}(\mathbf{x}) = \mathbf{W}_{\text{out}} \sigma(\mathbf{W}_{\text{in}} \mathbf{x} + \beta_{\text{in}}) + \beta_{\text{out}}, \quad (1)$$

where $\mathbf{W}_{\text{out}}^T, \mathbf{W}_{\text{in}} \in \mathbb{R}^{d_{\text{mlp}} \times d_{\text{model}}}$ are learned weight matrices, β_{in} and β_{out} are learned biases. The function σ denotes an element-wise nonlinear activation function, typically a GeLU [32]. In this paper, the term *neuron* refers to an entry in the MLP hidden state. We denote specific neurons in the model using the format `<layer>.<index>`. We use $\mathbf{n}$ to refer to the *activation values* of the neurons (i.e., the output of the activation function σ). Furthermore, we use $\mathbf{w}_{\text{out}}^{(i)} \in \mathbb{R}^{d_{\text{model}}}$ to indicate the output weights of neuron i (i.e., the i -th column of $\mathbf{W}_{\text{out}}$).

LayerNorm. Layer normalization (LayerNorm) is a commonly employed technique to improve the stability of the training process in deep neural networks [4]. Given a hidden state $\mathbf{x} \in \mathbb{R}^{d_{\text{model}}}$, LayerNorm applies the transformation

$$\text{LN}(\mathbf{x}) = \frac{\mathbf{x} - m(\mathbf{x})}{\sqrt{\text{Var}(\mathbf{x}) + \epsilon}} \odot \gamma + \beta. \quad (2)$$

In words, LayerNorm centers a hidden state by subtracting its mean ($m(\mathbf{x}) = \frac{1}{d_{\text{model}}} \sum_j \mathbf{x}_j$), re-scales it by its ϵ -adjusted standard deviation ($\sqrt{\text{Var}(\mathbf{x}) + \epsilon}$, where $\epsilon \in \mathbb{R}$), and applies an element-wise affine transformation determined by the learned parameters $\gamma, \beta \in \mathbb{R}^{d_{\text{model}}}$.

Unembedding. After the final transformer block, the final state of the residual stream is passed through a final LayerNorm and then projected onto the vocabulary space through an unembedding matrix $\mathbf{W}_{\text{U}} \in \mathbb{R}^{|\mathcal{V}| \times d_{\text{model}}}$. The logits obtained by mapping the residual stream final state onto the vocabulary space are then passed through a softmax function to convert them into a probability distribution $P_{\text{model}} \in \mathcal{V}$. Finally, the loss is computed using a function $\mathcal{L} : \mathcal{Y} \times \mathcal{V} \rightarrow \mathbb{R}_+$ (typically cross entropy for autoregressive language models). Following previous work [16], we employ standard weight pre-processing techniques [53]. In particular, we “fold” LayerNorm trainable parameters into the $\mathbf{W}_{\text{in}}$ and $\mathbf{W}_{\text{U}}$ matrices. We report the details about weight pre-processing in Appendix C.

3 Entropy neurons

Prior research examining individual neurons in GPT-2 language models [62] identified a category of neurons notable for their high norm and limited interaction with the unembedding matrix. Katz and Belinkov [42] speculate that these neurons may act as regularizers, contributing a constant bias to the residual stream. Gurnee et al. [28] independently rediscovered these neurons, termed them *entropy neurons*, and suggest that they play a role in regulating the entropy of the model’s output distribution. They show that interventions increasing the activation value of these neurons affect the entropy of the output distribution to a larger extent compared to other neurons. Despite these insights, the natural activation effects of these neurons on standard language inputs have not been examined, and *how* they carry out this entropy-modulating function remains unclear. In this section, we explore the mechanism of action of these neurons within the model.

3.1 Identifying entropy neurons

Our initial step is to identify which neurons are entropy neurons. We do this by searching for neurons with a high weight norm and a minimal impact on the logits. To detect minimal effect on the logits, we follow the heuristic used by Gurnee et al. [28], analyzing the variance in the effect of neurons on the logits. In particular, we project the last-layer neuron weights in GPT-2 Small onto the vocabulary space. This projection, sometimes referred to as logit attribution, represents an approximation of the neuron’s effect on the final prediction logits [55, 24, 15]. Then, we compute the variance of the normalized projection. That is, given a neuron with output weights $\mathbf{w}_{\text{out}}$, we compute

$$\text{LogitVar}(\mathbf{w}_{\text{out}}) = \text{Var} \left(\frac{\mathbf{W}_{\text{U}} \mathbf{w}_{\text{out}}}{\|\mathbf{W}_{\text{U}}\|_{\text{dim}=1} \|\mathbf{w}_{\text{out}}\|} \right), \quad (3)$$

⁴In particular, our study focuses on the final-layer MLPs due to their direct path to the final LayerNorm and unembedding, which excludes possible interactions with other transformer layers.




Figure 2: Identifying and Analyzing Entropy Neurons. (a) Neurons in GPT-2 Small displayed by their weight norm and variance in logit attribution. Entropy neurons (red) have high norm and low logit variance. (b) Causal graph showing the total effect and direct effect (bypassing LayerNorm) of a neuron on the model’s output. (c) Comparison of total and direct effects on model loss for entropy neurons and randomly selected neurons. (d) Singular values and cosine similarity between neuron output weights and singular vectors of $\mathbf{W}_U$. (e) Entropy neurons (red) show significant LayerNorm-mediated effects and high projection onto the null space (ρ). (f) Relationship between ρ and the LayerNorm-mediated effect in LLaMA2 7B. ρ is computed with $k = 40 \approx 0.01 * d_{\text{model}}$. Color represents absolute change in entropy upon ablation (ΔH).

where $\|\cdot\|_{\text{dim}=1}$ indicates column-wise norm.

This measure allows us to quantify how specific is the direct effect of a neuron on the output logits, with the underlying intuition that a diffused contribution (i.e., close to a constant added to all logits before applying the softmax) results in a small effect on the output probabilities. In the last layer of GPT-2 Small, we identify a subset of neurons with particularly low variance but substantial norm (Figure 2a). Since the optimization of the model’s weights preserves these high-norm low-composition neurons (that are expensive in terms of weight decay penalty [48]), it is likely that these neurons are important and exert their effect through a different mechanism than directly modifying the logits.

3.2 Mechanism of action

Given the low composition with the unembedding matrix, we hypothesize that most of the effect of entropy neurons on the model’s output is mediated by the re-scaling of the residual stream performed by the final LayerNorm. We test this hypothesis using a causal mediation analysis formulation [60, 64], illustrated in Figure 2b. We distinguish the *total* effect that a neuron has on the model’s output (represented by the two causal paths stemming from the neuron node in Figure 2b) from its *direct* effect (the green arrow in Figure 2b), which is not mediated by the change in the LayerNorm scale. Our hypothesis posits that the difference between the former and the latter is significantly larger for entropy neurons than for normal neurons.

We carry out an ablation experiment where we intervene on the activation value of a specific neuron by fixing it to its mean value across a reference distribution while constraining the LayerNorm scaling coefficient to remain constant. More formally, consider a neuron with index $i \in \{1, 2, \dots, d_{\text{mlp}}\}$ and denote by $\mathbf{n}_i \in \mathbb{R}$ its activation value. Given an input $x \in \mathcal{X}$, denote by $\mathbf{x}$ the last hidden state in the model (i.e., the output of the last transformer block), and denote by $\mathbf{x}^{-i}$ the last hidden state in the model obtained after mean-ablating neuron i :

$$\mathbf{x}^{-i} = \mathbf{x} + (\bar{\mathbf{n}}_i - \mathbf{n}_i) \mathbf{w}_{\text{out}}^{(i)}, \quad (4)$$

where $\bar{\mathbf{n}}_i$ is the mean activation value computed over a subset of $\mathcal{X}$.

We quantify the total effect of neuron i upon mean ablation by measuring the absolute variation in the model's loss $\mathcal{L}$,⁵ specifically:

$$\text{TE}(i) = \mathbb{E}_x \left[\left| \mathcal{L}(\mathbf{W}_U \text{LN}(\mathbf{x}), x) - \mathcal{L}(\mathbf{W}_U \text{LN}(\mathbf{x}^{-i}), x) \right| \right], \quad (5)$$

where the expectation is taken over a uniformly sampled subset of a corpus $\mathcal{X}$. Similarly, we quantify the direct effect of neuron i by preventing the LayerNorm denominator from varying:

$$\text{DE}_{\text{LN}}(i) = \mathbb{E}_x \left[\left| \mathcal{L}(\mathbf{W}_U \text{LN}(\mathbf{x}), x) - \mathcal{L}(\mathbf{W}_U (\frac{\mathbf{x}^{-i} - m(\mathbf{x}^{-i})}{\sqrt{\text{Var}(\mathbf{x}) + \epsilon}} \odot \gamma + \beta), x) \right| \right], \quad (6)$$

For each neuron, we measure the total and direct effects when its activation value is set to the mean across a dataset of 25600 tokens from the C4 Corpus [63] (additional experimental details are provided in Appendix D). We compare these metrics for six selected entropy neurons against those from 100 randomly selected neurons. The results (Figure 2c) show that in the randomly selected neurons there is no significant difference between their total and direct effects. On the other hand, the difference between these two quantities is substantial for entropy neurons, and in some cases the direct effect represents only a small fraction of the total effect.

These findings represent an interesting and novel example of how language models can use LayerNorm to indirectly manipulate the logit values. Such a mechanism could easily be overlooked by conventional analyses, which typically focus on direct logit attribution and fail to account for the normalization effects imposed by LayerNorm [16, 24, 46].

3.3 The unembedding has an effective null space

The unembedding is a linear map to a significantly higher-dimensional space (e.g., from 768 to 50,257 dimensions in GPT-2). Therefore, it is surprising that the output of a high-norm neuron would have little effect on the logits. We hypothesize that there is a subspace of the residual stream with an unusually low impact on the output that entropy neurons write onto. To investigate this, we compute the singular value decomposition (SVD) of the unembedding matrix $\mathbf{W}_U = \mathbf{U}\Sigma\mathbf{V}^T$.

Analyzing the singular values in Σ (thick blue line in Figure 2d), we observe that the bottom values are extremely small. In particular, we notice a sharp drop to values close to 0 around index 755. This observation indicates a remarkable phenomenon within the model: the training procedure optimizes the weights to preserve what effectively functions as a null space within $\mathbf{W}_U$. This finding is striking, as it suggests that the model deliberately limits its representational power by projecting a set of residual stream directions onto a small neighborhood of the 0-vector.

Next, we study the directions within the residual stream that entropy neurons write to. We do this by computing the cosine similarity between the output weights ($\mathbf{w}_{\text{out}}$) for each neuron and singular vectors $\mathbf{V} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$. The results, illustrated by the colored lines in Figure 2d, show that all the entropy neurons write almost exclusively to directions within the effective null space. These findings illustrate how entropy neurons leverage this null space to add norm to the residual stream without directly affecting the logit values.

3.4 Universality across model families

The relationship observed between entropy neurons and the unembedding null space suggests that entropy neurons may be identified by analyzing the proportion of their norm that is projected onto such a null space. To test this, given a neuron i , we compute the fraction ρ_i of the neuron's norm projected onto the effective null space $\mathbf{V}_0^T$. In particular, we define the effective null space of $\mathbf{W}_U$ as the space spanned by its bottom k singular vectors, i.e., $\mathbf{V}_0^T := \langle \mathbf{V}_{:,d_{\text{model}}-k}^T, \dots, \mathbf{V}_{:,d_{\text{model}}}^T \rangle$, and we compute $\rho_i = \|\mathbf{V}_0^T \mathbf{w}_{\text{out}}^{(i)}\| / \|\mathbf{w}_{\text{out}}^{(i)}\|$. Then, we measure the quantity $1 - \text{DE}(i)/\text{TE}(i)$, which represents the fraction of the neuron's effect that is mediated by the final LayerNorm.⁶ We

⁵For readability, here we omit the final softmax to normalize the logits before computing the loss.

⁶We opt for representing the LayerNorm-mediated effect by $1 - \text{DE}(i)/\text{TE}(i)$, rather than computing the indirect effect of the LayerNorm scale because a large indirect effect might highlight neurons with high norm and large direct effect, such as a neuron that boosts a specific logit. In our setting, we are interested in neurons that actually have low direct effect, which is not implied by a large indirect effect.




Figure 3: **Token Frequency Neurons in Pythia 410M.** (a) $D_{KL}(P_{\text{freq}} || P_{\text{model}})$ and Entropy are correlated negatively. (b) Scatter plot of neurons highlighting token frequency neurons (in green), with high effect on $D_{KL}(P_{\text{freq}} || P_{\text{model}})$, significantly mediated by the token frequency direction. (c) Box plots showing substantial difference in total vs. direct effect in token frequency neurons.

compare these two quantities for each neuron at the last layer in GPT-2 Small (Figure 2e), with $k = 12$. The results show that neurons with the largest projections onto the null space have the most significant LayerNorm-mediated effects on the model’s output. These findings connect with recent work highlighting the importance of the bottom $\mathbf{W}_U$ singular vectors [10] and represent further evidence supporting our hypothesis about the entropy neurons’ mechanism of action.

We extend our investigation of entropy neurons to models beyond the GPT-2 family. In Figure 2f, we analyze the proportion of norm projected onto the bottom singular vectors of $\mathbf{W}_U$ (i.e., ρ) and the influence of LayerNorm on the output for neurons in the last layer of LLaMA2 7B [70].⁷ Similar to our observations with GPT-2 Small, the results indicate the existence of a distinct group of neurons that predominantly write onto the effective null space of $\mathbf{W}_U$. These neurons significantly influence the model’s output via LayerNorm, confirming the presence of entropy neurons across different model families. As entropy neurons act by re-scaling the residual stream and have no effect that is specific to a subset of the tokens, their activation has a large effect on the entropy of the output distribution but little impact on the relative ranking of the tokens. That is, they affect the model’s confidence in the prediction without affecting the actual token being predicted. This is illustrated in Figure 1. We repeat these analyses on GPT-2 Medium, Pythia 410M and 1B[6], Phi 2 [38], and Gemma 2B [68], obtaining overall consistent results (Appendix F).

4 Token frequency neurons

The token frequency distribution (i.e., the distribution of unigrams over a corpus) can serve as a reliable baseline for next-token prediction, especially in scenarios of high uncertainty [11, 51]. Thus, we hypothesize that a useful general-purpose mechanism for regulating confidence in language models could involve modulating the distance between the output distribution and the token frequency distribution. We explore this hypothesis using the 410M-parameter Pythia model, for which the training corpus has been publicly released (The Pile; 20). For this corpus, we compute the empirical token frequency distribution P_{freq} . We observe that the entropy of the model’s output distribution is in fact negatively correlated with its Kullback-Leibler (KL) divergence from the empirical token frequency distribution: as the confidence of the model decreases (higher entropy), its output distribution tends to be closer to the empirical token frequency distribution (Figure 3a).

As for entropy neurons, we aim to identify neurons whose operational mechanism relies on a particular subspace—in this case, the direction corresponding to the token frequency distribution. We obtain such a direction by computing the logit vector $\mathbf{v}_{\text{freq}} \in \mathbb{R}^{|\mathcal{V}|}$ by centering the log-probability values for each token in $\mathcal{V}$. That is, $\mathbf{v}_{\text{freq},i} = \log(\mathbf{p}_{\text{freq},i}) - m(\log(\mathbf{p}_{\text{freq}}))$, where $\mathbf{p}_{\text{freq}} \in \mathbb{R}^{|\mathcal{V}|}$ the vector of the token frequency probabilities values, and $m : \mathbb{R}^{|\mathcal{V}|} \rightarrow \mathbb{R}$ indicates the mean function.

To identify the neurons that rely on this direction to modulate the model’s output, we conduct an ablation experiment similar to the one described in §3.2. In this experiment, we mean-ablate neurons and assess the absolute change in model loss while preserving the value of the logits along the

⁷LLaMA actually uses RMSNorm [80] which differs from LayerNorm in the absence of re-centering and bias term. However, our experimental procedure is not affected by this difference.




Figure 4: **Examples of Neuron Activity in Language Models.** (a) Change in loss after ablation of entropy neuron 11.2378 in GPT-2 Small. Color indicates reciprocal rank (RR) of the correct token prediction. (b) Activation of neuron 11.2378 on an example from the C4 Corpus. The neuron mitigates a loss spike at the token “Mes,” after which the model predicts “otherapy.” (c) Change in entropy and KL divergence on correct tokens (RR = 1) post ablation of neuron 23.417 in Pythia 410M. The neuron increases entropy and aligns the model’s output with the token frequency distribution.

direction of $\mathbf{v}_{\text{freq}}$ as constant. Denote by $\mathbf{l} = \mathbf{W}_U \text{LN}(\mathbf{x})$ the logits produced by the model on input $x \in \mathcal{X}$, and denote by $\mathbf{l}^{-i} = \mathbf{W}_U \text{LN}(\mathbf{x}^{-i})$ the value of the logits obtained on the same input after the ablation of a neuron i (where $\mathbf{x}^{-i}$ is defined in Eq. (4)). To measure the effect of neuron i that is not mediated by the token frequency direction (i.e., its *direct* effect), we obtain the adjusted value of the post-ablation logits $\tilde{\mathbf{l}}^{-i}$ as

$$\tilde{\mathbf{l}}^{-i} = \mathbf{l}^{-i} + \frac{(\mathbf{l}^T \mathbf{v}_{\text{freq}}) - (\mathbf{l}^{-iT} \mathbf{v}_{\text{freq}})}{\|\mathbf{v}_{\text{freq}}\|^2} \mathbf{v}_{\text{freq}}. \quad (7)$$

Then we compute the total and direct effects of neuron i as

$$\text{TE}(i) = \mathbb{E}_{x \sim \mathcal{X}} \left[\left| \mathcal{L}(\mathbf{l}, x) - \mathcal{L}(\mathbf{l}^{-i}, x) \right| \right], \text{ and } \text{DE}_{\text{freq}}(i) = \mathbb{E}_{x \sim \mathcal{X}} \left[\left| \mathcal{L}(\mathbf{l}, x) - \mathcal{L}(\tilde{\mathbf{l}}^{-i}, x) \right| \right]. \quad (8)$$

In Figure 3b, we report the results comparing the token frequency-mediated effect against the average absolute change in the KL divergence between P_{model} and P_{freq} for neurons at the final layer of Pythia 410M. We also highlight the entropy neurons with the highest LayerNorm-mediated effect (see Appendix F for comprehensive results on entropy neurons in Pythia 410M). We observe that there are neurons whose effect on P_{model} is substantially mediated by the token frequency direction (i.e., positive value along the x -axis in Figure 3b). These neurons, upon ablation, lead to significant variation in $D_{\text{KL}}(P_{\text{freq}} \| P_{\text{model}})$ (large positive value along the y -axis in Figure 3b), comparable to the variations caused by some entropy neurons.

These results validate the presence of components (which we term *token frequency neurons*) that affect the model’s output distribution by bringing it closer or away from the token frequency distribution by *directly* modulating it. Figure 3c focuses on 5 selected token frequency neurons and shows their impact on the loss and its decrease upon the inhibition of the token frequency direction (Eq. (7)), which parallels the LayerNorm mediation for entropy neurons discussed in §3. We additionally show the presence of token frequency neurons in GPT-2 Small and Pythia 1B in Appendix G.

5 Examples of neuron activity

We have shown that confidence-regulation neurons improve model performance by calibrating its output across the entire distribution. However, it remains unclear what this looks like in practice. To better illustrate the role that these neurons play in language models, we provide additional analyses about the changes in the model output induced by particular entropy and token frequency neurons.

Entropy neurons. In GPT-2 Small, we examine the change in entropy induced by one of the strongest entropy neurons identified in §3 (11.2378). To study this, we conduct an experiment in which we clip the activation value of the neuron to its mean and measure the resulting change in the model output.⁸ We analyze the difference in loss caused by the ablation of the neuron on

⁸“Clipping” refers to setting the activation value of the neuron to its mean only when it exceeds it, unlike mean ablation, which sets the activation value to its mean regardless.




Figure 5: Entropy Neurons on Induction. (a) Activations, entropy, and loss across duplicated 200-token input sequences. (b) The effect of clip mean-ablation of specific entropy neurons. Neuron 11.2378 shows the most significant impact, with up to a 70% reduction in entropy. (c) BOS ablation of induction heads: Upon the ablation of three induction heads in GPT-2 Small, the activation of entropy neuron 11.2378 decreases substantially.

individual tokens and compare it to the initial loss value. This analysis is depicted in Figure 4a, which also illustrates the reciprocal rank of the correct token prediction in the model’s output distribution (RR). These observations suggest that the entropy-increasing function of neuron 11.2378 acts as a hedging mechanism: it marginally increases the loss on correct model predictions (low initial loss) but prevents the loss from spiking when the model confidently predicts a wrong token (high initial loss). In Figure 4b, we provide an example of input, taken from the C4 corpus [63], on which the model confidently predicts a wrong token (“otherapy”), and neuron 11.2378 activates to mitigate the loss spike. We perform the same analysis for the strongest entropy neuron in LLaMA2 7B in Appendix F.1.

Token frequency neurons. We next focus on neuron 23.417 in Pythia 410M (identified in §4). This is a token frequency neuron: when its activation is increased, it pushes the model’s output P_{model} towards the token frequency distribution P_{freq} . This is illustrated by Figure 4c, showing the KL divergence between the two distributions pre- and post-ablation. As with increased entropy, bringing P_{model} closer to P_{freq} decreases the model’s confidence (positive values along the x -axis correlate with negative y values in Figure 4c). This leads to an increase in loss on correct predictions, reflected in the negative change in loss upon ablation (darker points in Figure 4c). Interestingly, we find that the projection of this neurons’ $\mathbf{w}_{\text{out}}$ onto the frequency direction $\mathbf{v}_{\text{freq}}$ is negative, suggesting that this neuron is suppressing common tokens and promoting rare ones. This indicates that the model is on average biased to predict common tokens even more frequently than their frequency would dictate. This way, lowering confidence by pushing P_{model} closer to P_{freq} requires suppressing the token frequency direction, which is the opposite of what one would have naively expected.

6 Case study: Induction

As a more detailed case study to examine the active role and confidence-regulating function of entropy neurons, we focus on the setting of induction. Induction occurs when there is a repeated subsequence of tokens in the model’s input, and the model must detect the earlier repeat and continue it. Mechanistically, this is implemented by specialized attention heads, called induction heads [16, 57], which attend from some token A to the token B that immediately follows A in the early occurrence and predict B (AB. . . A $\rightarrow$ B). Repeated text frequently occurs in natural language (e.g., someone’s name) and, during a repeated subsequence, the next token can often be predicted with very high confidence.

6.1 Entropy neurons

To analyze this phenomenon, we create input sequences by selecting 100 tokens from C4 [63] and duplicating them to form a 200-token input sequence. Across 100 such sequences, we measure GPT-2 Small’s performance and observe a significant decrease in both average loss and entropy during the second occurrence of the sequence (solid lines in Figure 5a). Additionally, we track the activation values of the six entropy neurons in GPT-2 Small analyzed in §3 (dotted lines in Figure 5a). For four of these neurons (11.584, 11.2378, 11.2123, and 11.2910), we note substantial change in average activation values, suggesting that they may play an important role. Furthermore, among final-layer neurons, entropy neurons exhibit the largest change in activation during induction (Appendix H.3, Figure 11).

To further explore the specific effects of entropy neurons on the model’s output entropy and loss, we conduct an experiment where, during the second occurrence of the sequence, we clip the activation values of the neurons to their mean values from the first occurrence. The results (Figure 5b) reveal that while the ablation of some neurons such as 11.584 and 11.2123 leads to a slight increase in output entropy, the most significant impact was observed with neuron 11.2378. This results in up to a 70% reduction in entropy, suggesting its role in boosting entropy to counterbalance the model’s high confidence during induction. We additionally study the effect of token frequency neurons during induction and we report the results in Appendix H.2, revealing that these neurons also lead to a significant change in entropy compared to randomly selected neurons. We repeat these analyses in naturally occurring induction cases identified in the C4 corpus, obtaining consistent results (Appendix H.1).

6.2 Induction head–entropy neuron interaction

In order to verify that induction heads have a causal effect on the activation of our entropy neurons, we perform *BOS ablations* [31]. That is, motivated by the observation that attention heads frequently attend to BOS when inactive (e.g., when induction is not occurring) [72], we set the attention pattern of an attention head to always attend to the first token of the sequence (i.e., BOS) and record the neuron’s activation. We perform this procedure on the top 3 induction heads, selected according to the prefix matching score introduced by Olsson et al. [57]. In GPT-2 Small, these heads are L5H1, L5H5, and L6H9, respectively. As baselines, we randomly select 6 heads in layers 5 and 6 that, like our induction heads, have mean attention to BOS > 0.6 but have prefix matching score < 0.1. For both entropy neurons and baselines, we ablate heads individually and in combinations of up to 3 heads.

Figure 5c shows that ablating L5H1 (ind. head 1) leads to a significant decrease in 11.2378’s mean activation. Ablating L5H5 (ind. head 2) or L6H9 (ind. head 3) individually does not make a substantial difference, but ablating them alongside L5H1 further decreases activation, bringing 11.2378 close to its activation on the first 100 tokens. We perform additional BOS ablations for other entropy neurons in Appendix H and further study the induction head–entropy neuron interaction in Appendix I, providing preliminary evidence that entropy neurons respond to an internal ‘induction-has-occurred’ signal that is produced by induction heads.

7 Related work

Uncertainty in language models. Uncertainty quantification for machine learning models has been extensively studied to assess the reliability of model predictions [19, 21]. With the increasing adoption of LLMs, accurately determining their confidence has drawn significant attention [22]. Existing research has explored various strategies to calibrate model output, including ensembling model predictions [76, 34], fine-tuning [39, 40], and prompting models to explicitly express their confidence [47, 40, 65, 69]. Additionally, efforts have been made to evaluate the truthfulness [9, 3, 50] and uncertainty [35] of model outputs by inspecting internal representations. However, the investigation of internal, general-purpose mechanisms determining model confidence remains underexplored.

Interpretability of LLMs’ components. The analysis of language models’ internal components and mechanisms is an increasingly popular area of research (we refer to Ferrando et al. [18] for an overview). Prior work has explored how to attribute and localize model behavior [73, 23, 52], and to identify specific algorithms that models implement to perform tasks [54, 75, 29, 67, 12]. Additionally, many studies have focused on investigating individual neurons [41, 61, 14, 2, 13, 74, 27, 7]. However, neurons are not always the correct unit of analysis, as they can be polysemantic [56, 17]. Recent work has employed sparse autoencoders (SAEs) to find more interpretable linear combinations of neurons [79, 8, 36]. Despite this, in the context of entropy modulation, prior findings of relevant neurons motivated us to focus on neurons.

Closest related work. An anomalous class of neurons characterized by high norm and low composition with the unembedding matrix was observed by Katz and Belinkov [42], though this was not linked to the regulation of output distribution entropy. This class of neurons was independently rediscovered by Gurnee et al. [28] while investigating the universality of neurons in GPT-2 models. Our work connects the concept of entropy neurons to the presence of an effective null space represented by the bottom singular vectors of $\mathbf{W}_U$. Our findings highlight the significance of this seemingly unimportant subspace within the residual stream and align with the observations of Cancedda [10], who associated the bottom $\mathbf{W}_U$ singular vectors with the phenomenon of the attention sink [78].

8 Conclusion

This paper presents an analysis of the mechanisms by which LLMs manage and express uncertainty, focusing on two specific components: entropy neurons and token frequency neurons. We show that entropy neurons, with their high weight norm and low direct interaction with the unembedding matrix, affect the model’s output via the final LayerNorm. We also introduce and investigate token frequency neurons, which adjust the model’s output by aligning it closer to or further from the token frequency distribution. In a case study on induction, we demonstrate the practical implications of these neurons. We show that one example role of entropy neurons is acting as a hedging mechanism to manage confidence in repetitive sequence scenarios. Some limitations of our work include focusing only on two types of components in the neuron basis, relying on proxies for confidence, and observing varying effects across models. We thoroughly discuss the limitations of our work in Appendix A. This study represents the first thorough investigation into the mechanisms that LLMs might use for confidence calibration, providing insights that can guide further research in this area.

Acknowledgments

This work was conducted as part of the ML Alignment & Theory Scholars (MATS) Program. Throughout the project, we received valuable input from fellow MATS participants. In particular, we would like to thank Andy Ardit, Joseph Bloom, Connor Kissane, and Robert Krzyzanowski. We also extend our gratitude to McKenna Fitzgerald for organizational support. Additionally, we would like to thank Arthur Conmy, Giulia Lanzillotta, Mazda Moayeri, Kumar Shridhar, and Vilém Zouhar for useful discussions and comments on our work. Yonatan was supported by the Israel Science Foundation (grant no. 448/20), an Azrieli Foundation Early Career Faculty Fellowship, and an AI Alignment grant from Open Philanthropy. Alessandro and Ben were supported by a Long-term Future Fund. Alessandro is further supported by armasuisse Science and Technology through a CYD Doctoral Fellowship. Ben is supported through an EPSRC Doctoral Training Partnership Grant. The project was also supported by UK’s innovation agency (Innovate UK) grant 10098112 (project name ASIMOV: AI-as-a-Service) and by the European Union under the Horizon Europe vera.ai project (grant 101070093).

Author Contribution

Alessandro and Ben led the project, carried out the experiments, and wrote the paper. Alessandro proposed the causal formulation and conducted most of the experiments validating the entropy neurons’ mechanism of action. He also performed experiments to identify entropy neurons and token frequency neurons across multiple models and wrote the majority of §§ 1 to 5. Ben carried out numerous validation experiments to ensure the robustness of our findings and obtained preliminary results on the universality of entropy neurons. Ben also led the experiments for the induction case study and wrote most of §6 and Apps. H and I. Wes gave frequent and detailed feedback on experiment design and paper write-up. Yonatan, Xingyi, and Mrinmaya provided valuable input on experiment design and presentation of the results. Neel supervised the project and provided guidance at all stages on experiment design and analysis, in addition to editing the paper.

References

- [1] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016. URL <https://arxiv.org/abs/1606.06565>. 1
- [2] Omer Antverg and Yonatan Belinkov. On the pitfalls of analyzing individual neurons in language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=8uz0EWPQIMu>. 9
- [3] Amos Azaria and Tom Mitchell. The internal state of an LLM knows when it’s lying. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, Singapore, dec 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.68. URL <https://aclanthology.org/2023.findings-emnlp.68>. 9

- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. URL <https://arxiv.org/abs/1607.06450>. 2, 3
- [5] Yoshua Bengio, Geoffrey Hinton, Andrew Yao, Dawn Song, Pieter Abbeel, Trevor Darrell, Yuval Noah Harari, Ya-Qin Zhang, Lan Xue, Shai Shalev-Shwartz, Gillian Hadfield, Jeff Clune, Tegan Maharaj, Frank Hutter, Atılım Güneş Baydin, Sheila McIlraith, Qiqi Gao, Ashwin Acharya, David Krueger, Anca Dragan, Philip Torr, Stuart Russell, Daniel Kahneman, Jan Brauner, and Sören Mindermann. Managing extreme AI risks amid rapid progress. *Science*, 384 (6698):842–845, 2024. doi: 10.1126/science.adn0117. URL <https://www.science.org/doi/abs/10.1126/science.adn0117>. 1
- [6] 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. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR, 2023. URL <https://proceedings.mlr.press/v202/biderman23a.html>. 2, 6, 19
- [7] Steven Bills, Nick Cammarata, Dan Mossing, Henk Tillman, Leo Gao, Gabriel Goh, Ilya Sutskever, Jan Leike, Jeff Wu, and William Saunders. Language models can explain neurons in language models. *OpenAI Blog*, 2023. URL <https://openai.com/index/language-models-can-explain-neurons-in-language-models/>. 9
- [8] Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023. URL <https://transformer-circuits.pub/2023/monosemantic-features/index.html>. 9
- [9] Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language models without supervision. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=ETKGuby0hcs>. 9
- [10] Nicola Cancedda. Spectral filters, dark signals, and attention sinks. *arXiv preprint arXiv:2402.09221*, 2024. URL <https://arxiv.org/abs/2402.09221>. 6, 9
- [11] Tyler A. Chang and Benjamin K. Bergen. Word acquisition in neural language models. *Transactions of the Association for Computational Linguistics*, 10:1–16, 01 2022. ISSN 2307-387X. doi: 10.1162/tacl_a_00444. URL https://doi.org/10.1162/tacl_a_00444. 2, 6
- [12] Arthur Conmy, Augustine Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. Towards automated circuit discovery for mechanistic interpretability. *Advances in Neural Information Processing Systems*, 36:16318–16352, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/34e1dbe95d34d7ebaf99b9bcaeb5b2be-Abstract-Conference.html. 9
- [13] Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. Knowledge neurons in pretrained transformers. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland, may 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.581. URL <https://aclanthology.org/2022.acl-long.581>. 9
- [14] Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, Anthony Bau, and James Glass. What is one grain of sand in the desert? analyzing individual neurons in deep nlp models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6309–6317, 2019. URL <https://aaai.org/ojs/index.php/AAAI/article/view/4592>. 9
- [15] Guy Dar, Mor Geva, Ankit Gupta, and Jonathan Berant. Analyzing transformers in embedding space. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,

pages 16124–16170, Toronto, Canada, jul 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.893. URL <https://aclanthology.org/2023.acl-long.893>. 3

- [16] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. URL <https://transformer-circuits.pub/2021/framework/index.html>. 2, 3, 5, 8
- [17] Nelson Elhage, Tristan Hume, Catherine Olsson, Neel Nanda, Tom Henighan, Scott Johnston, Sheer ElShowk, Nicholas Joseph, Nova DasSarma, Ben Mann, Danny Hernandez, Amanda Askell, Kamal Ndousse, Andy Jones, Dawn Drain, Anna Chen, Yuntao Bai, Deep Ganguli, Liane Lovitt, Zac Hatfield-Dodds, Jackson Kernion, Tom Conerly, Shauna Kravec, Stanislav Fort, Saurav Kadavath, Josh Jacobson, Eli Tran-Johnson, Jared Kaplan, Jack Clark, Tom Brown, Sam McCandlish, Dario Amodei, and Christopher Olah. Softmax linear units. *Transformer Circuits Thread*, 2022. URL <https://transformer-circuits.pub/2022/solu/index.html>. 9
- [18] Javier Ferrando, Gabriele Sarti, Arianna Bisazza, and Marta R Costa-jussà. A primer on the inner workings of transformer-based language models. *arXiv preprint arXiv:2405.00208*, 2024. URL <https://arxiv.org/abs/2405.00208>. 9
- [19] Yarin Gal et al. Uncertainty in deep learning, 2016. URL <http://106.54.215.74/2019/20190729-liuzy.pdf>. 1, 9
- [20] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020. URL <https://arxiv.org/abs/2101.00027>. 6, 19
- [21] Jakob Gawlikowski, Cedrique Rovile Njieutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56 (Suppl 1):1513–1589, 2023. URL <https://link.springer.com/article/10.1007/s10462-023-10562-9>. 1, 9
- [22] Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koeppl, Preslav Nakov, and Iryna Gurevych. A survey of language model confidence estimation and calibration. *arXiv preprint arXiv:2311.08298*, 2023. URL <https://arxiv.org/abs/2311.08298>. 1, 9
- [23] Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. Transformer feed-forward layers are key-value memories. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495, Online and Punta Cana, Dominican Republic, nov 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.446. URL <https://aclanthology.org/2021.emnlp-main.446>. 9
- [24] Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45, Abu Dhabi, United Arab Emirates, dec 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.3. URL <https://aclanthology.org/2022.emnlp-main.3>. 2, 3, 5
- [25] Aaron Gokaslan and Vanya Cohen. Openwebtext corpus, 2019. URL <http://SkyLion007.github.io/OpenWebTextCorpus>. 21
- [26] Gemini Team Google. Gemini: A family of highly capable multimodal models, 2024. URL <https://arxiv.org/abs/2312.11805>. 2

- [27] Wes Gurnee, Neel Nanda, Matthew Pauly, Katherine Harvey, Dmitrii Troitskii, and Dimitris Bertsimas. Finding neurons in a haystack: Case studies with sparse probing. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=JYs1R9IMJr>. 9
- [28] Wes Gurnee, Theo Horsley, Zifan Carl Guo, Tara Rezaei Kheirkhah, Qinyi Sun, Will Hathaway, Neel Nanda, and Dimitris Bertsimas. Universal neurons in GPT2 language models. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=ZeI104QZ8I>. 1, 2, 3, 9
- [29] Michael Hanna, Ollie Liu, and Alexandre Variengien. How does GPT-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=p4PckNQR8k>. 9
- [30] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Hal-dane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, sep 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>. 19
- [31] Stefan Heimersheim and Jett Janiak. A circuit for python docstrings in a 4-layer attention-only transformer. Alignment Forum, 2023. URL <https://www.alignmentforum.org/posts/u6KXXmKFbXfWzoAXn/a-circuit-for-python-docstrings-in-a-4-layer-attention-only>. 9
- [32] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. URL <https://arxiv.org/abs/1606.08415>. 3
- [33] Dan Hendrycks, Nicholas Carlini, John Schulman, and Jacob Steinhardt. Unsolved problems in ml safety. *arXiv preprint arXiv:2109.13916*, 2021. URL <https://arxiv.org/abs/2109.13916>. 1
- [34] Bairu Hou, Yujian Liu, Kaizhi Qian, Jacob Andreas, Shiyu Chang, and Yang Zhang. Decomposing uncertainty for large language models through input clarification ensembling. *arXiv preprint arXiv:2311.08718*, 2023. URL <https://arxiv.org/abs/2311.08718>. 9
- [35] Brian RY Huang and Joe Kwon. Does it know?: Probing for uncertainty in language model latent beliefs. In *NeurIPS Workshop on Attributing Model Behavior at Scale*, 2023. URL <https://openreview.net/forum?id=uSvN2oozRK>. 9
- [36] Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=F76bwRSLek>. 9
- [37] Plotly Technologies Inc. Collaborative data science, 2015. URL <https://plot.ly>. 19
- [38] Mojan Javaheripi, Sébastien Bubeck, Marah Abdin, Jyoti Aneja, Sebastien Bubeck, Caio César Teodoro Mendes, Weizhu Chen, Allie Del Giorno, Ronen Eldan, Sivakanth Gopi, et al. Phi-2: The surprising power of small language models. *Microsoft Research Blog*, 2023. URL <https://www.microsoft.com/en-us/research/blog/phi-2-the-surprising-power-of-small-language-models/>. 2, 6
- [39] Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977, 2021. doi: 10.1162/tac1_a_00407. URL <https://aclanthology.org/2021.tac1-1.57>. 9

- [40] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022. URL <https://arxiv.org/abs/2207.05221>. 1, 9
- [41] Andrej Karpathy, Justin Johnson, and Li Fei-Fei. Visualizing and understanding recurrent networks. *arXiv preprint arXiv:1506.02078*, 2015. URL <https://arxiv.org/abs/1506.02078>. 9
- [42] Shahar Katz and Yonatan Belinkov. VISIT: Visualizing and interpreting the semantic information flow of transformers. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14094–14113, Singapore, dec 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.939. URL <https://aclanthology.org/2023.findings-emnlp.939>. 1, 3, 9
- [43] Connor Kissane, Robert Krzyzanowski, Arthur Conmy, and Neel Nanda. Sparse autoencoders work on attention layer outputs. Alignment Forum, 2024. URL <https://www.alignmentforum.org/posts/DtdzGwFh9dCfsekZZ>. 25
- [44] Robert Krzyzanowski, Connor Kissane, Arthur Conmy, and Neel Nanda. We inspected every head in GPT-2 small using SAEs so you don’t have to. Alignment Forum, 2024. URL <https://www.alignmentforum.org/posts/xmegew5mqiBsvoaim/we-inspected-every-head-in-gpt-2-small-using-saes-so-you-don>. 23
- [45] Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=VD-AYtP0dve>. 1
- [46] Tom Lieberum, Matthew Rahtz, János Kramár, Geoffrey Irving, Rohin Shah, and Vladimir Mikulik. Does circuit analysis interpretability scale? evidence from multiple choice capabilities in chinchilla. *arXiv preprint arXiv:2307.09458*, 2023. URL <https://arxiv.org/abs/2307.09458>. 2, 5
- [47] Stephanie Lin, Jacob Hilton, and Owain Evans. Teaching models to express their uncertainty in words. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=8s8K2UZGTZ>. 1, 9
- [48] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam, 2018. URL <https://openreview.net/forum?id=rk6qdGgCZ>. 2, 4
- [49] Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=jN5y-zb5Q7m>. 1
- [50] Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets, 2024. URL <https://openreview.net/forum?id=CeJEfNKstt>. 9
- [51] Clara Meister, Wojciech Stokowiec, Tiago Pimentel, Lei Yu, Laura Rimell, and Adhiguna Kuncoro. A natural bias for language generation models. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 243–255, Toronto, Canada, jul 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-short.22. URL <https://aclanthology.org/2023.acl-short.22>. 2, 6
- [52] Kevin Meng, David Bau, Alex J Andonian, and Yonatan Belinkov. Locating and editing factual associations in GPT. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=-h6WAS6eE4>. 9
- [53] Neel Nanda and Joseph Bloom. Transformerlens, 2022. URL <https://github.com/TransformerLensOrg/TransformerLens>. 3, 18, 19, 24

- [54] Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=9XFSbDPmdW>. 9
- [55] Nostalgebraist. Interpreting gpt: The logit lens, 2020. URL <https://www.alignmentforum.org/posts/AcKRB8wDpdAN6v6ru/interpreting-gpt-the-logit-lens>. 3
- [56] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2017. doi: 10.23915/distill.00007. URL <https://distill.pub/2017/feature-visualization>. 9
- [57] Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. URL <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>. 2, 8, 9
- [58] OpenAI. GPT-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>. 1, 2
- [59] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf. 19
- [60] Judea Pearl. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, UAI'01, page 411–420, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558608001. URL <https://dl.acm.org/doi/10.5555/2074022.2074073>. 4
- [61] Alec Radford, Rafal Jozefowicz, and Ilya Sutskever. Learning to generate reviews and discovering sentiment, 2018. URL <https://openreview.net/forum?id=SJ71VXZAZ>. 9
- [62] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. URL <https://insightcivic.s3.us-east-1.amazonaws.com/language-models.pdf>. 2, 3
- [63] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. URL <https://www.jmlr.org/papers/v21/20-074.html>. 5, 8, 19
- [64] James M Robins and Sander Greenland. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(2):143–155, 1992. URL https://journals.lww.com/epidem/abstract/1992/03000/identifiability_and_exchangeability_for_direct_and.13.aspx. 4
- [65] Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Lee Boyd-Graber, and Lijuan Wang. Prompting GPT-3 to be reliable. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=98p5x51L5af>. 1, 9
- [66] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. URL https://www.jmlr.org/papers/volume15/srivastava14a/srivastava14a.pdf?utm_content=buffer79b43&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer. 19

- [67] Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7035–7052, Singapore, dec 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.435. URL <https://aclanthology.org/2023.emnlp-main.435>. 9
- [68] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024. URL <https://arxiv.org/abs/2403.08295>. 2, 6
- [69] Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore, dec 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.330. URL <https://aclanthology.org/2023.emnlp-main.330>. 9
- [70] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. URL <https://arxiv.org/abs/2307.09288>. 2, 6
- [71] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. URL <https://proceedings.neurips.cc/paper/7181-attention-is-all>. 2
- [72] Jesse Vig and Yonatan Belinkov. Analyzing the structure of attention in a transformer language model. In Tal Linzen, Grzegorz Chrupała, Yonatan Belinkov, and Dieuwke Hupkes, editors, *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 63–76, Florence, Italy, aug 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4808. URL <https://aclanthology.org/W19-4808>. 9
- [73] Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. Investigating gender bias in language models using causal mediation analysis. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12388–12401. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/92650b2e92217715fe312e6fa7b90d82-Paper.pdf. 9
- [74] Elena Voita, Javier Ferrando, and Christoforos Nalmpantis. Neurons in large language models: Dead, n-gram, positional. *arXiv preprint arXiv:2309.04827*, 2023. URL <https://arxiv.org/abs/2309.04827>. 9
- [75] Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: A circuit for indirect object identification in GPT-2 small. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=NpsVSN6o4u1>. 9
- [76] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=1PL1NIMMrw>. 9
- [77] Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56 – 61, 2010. doi: 10.25080/Majora-92bf1922-00a. URL <https://scipy.org>. 19

- [78] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. In *The Twelfth International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=NG7sS51zVF>. 9
- [79] Zeyu Yun, Yubei Chen, Bruno Olshausen, and Yann LeCun. Transformer visualization via dictionary learning: Contextualized embedding as a linear superposition of transformer factors. In Eneko Agirre, Marianna Apidianaki, and Ivan Vulić, editors, *Proceedings of Deep Learning Inside Out (DeeLIO): The 2nd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 1–10, Online, jun 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.deelio-1.1. URL <https://aclanthology.org/2021.deelio-1.1.9>
- [80] Biao Zhang and Rico Sennrich. Root mean square layer normalization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/1e8a19426224ca89e83cef47f1e7f53b-Paper.pdf. 6
- [81] Vilém Zouhar. Ryanize bib, 2023. URL <https://github.com/zouharvi/ryanize-bib>. 19

A Limitations

Our study focuses exclusively on two particular types of final-layer components: entropy neurons and token frequency neurons. However, it is likely that there are multiple other components and mechanisms within language models that contribute to confidence regulation. Exploring these additional components could provide a more comprehensive understanding of how language models manage uncertainty.

We also acknowledge the lack of a precise definition of confidence. Instead, we rely on rough proxies such as entropy and the distance from the token frequency distribution. These measures, while useful, may not fully capture the concept of confidence in language models.

Although we observed the presence of entropy neurons in most of the models we studied, the fraction of their effect explained by LayerNorm re-scaling varies. For instance, in Pythia and Gemma, this fraction is around 40%, which is significantly less than the 80% observed in GPT-2, LLaMA2 7B, and Phi-2. We believe that variations in model architecture and training hyper-parameters may explain these inconsistencies. Although we do not explore these factors in detail, we discovered that the application of dropout during training leads to the emergence of a larger null space (Appendix E), which we hypothesize may encourage the formation of entropy neurons. Future research should rigorously examine the factors that determine the emergence of entropy neurons, as well as other mechanisms that regulate model confidence.

Lastly, our study is primarily focused on how entropy and token frequency neurons contribute to confidence regulation in next-token prediction. An interesting direction for future research would be to investigate confidence-regulating neurons in broader contexts and real-world tasks, such as question-answering or reasoning. This could provide deeper insights into the practical applications and limitations of these neurons in diverse scenarios.

B Broader Impacts

This research enhances the understanding of how LLMs regulate confidence in their predictions through entropy neurons and token frequency neurons. While our study primarily advances our technical understanding of language models, some broader impacts should be considered. Enhanced confidence calibration might lead to biased or discriminatory decisions if not adequately mitigated, perpetuating or amplifying existing biases in training data. Additionally, insights from this research could facilitate extraction attacks, compromising data privacy, and be exploited for adversarial attacks, undermining model security. To mitigate these risks, it is essential to integrate robust fairness checks and bias mitigation strategies during model training and deployment. Implementing privacy-preserving techniques, such as differential privacy, can help protect against data leakage. Additionally, developing adversarial defense mechanisms and deploying real-time monitoring systems can counteract potential misuse.

C Weight Pre-processing

With the purpose of removing irrelevant components and other parameterization degrees of freedom, we employ a set of standard weight pre-processing techniques [53]. This way, cosine similarities and other weight computations have a mean of 0.

Folding in Layer Norm. Most layer norm implementations include trainable parameters $\gamma \in \mathbb{R}^n$ and $\beta \in \mathbb{R}^n$ (see Eq. (2)). To account for these, we “fold” the layer norm parameters into $\mathbf{W}_{\text{in}}$ by treating the layer norm parameters as equivalent to a linear layer and then combining the adjacent linear layers. We create effective weights as follows:

$$\mathbf{W}_{\text{eff}} = \mathbf{W}_{\text{in}} \cdot \text{diag}(\gamma), \quad \beta_{\text{eff}} = \beta_{\text{in}} + \mathbf{W}_{\text{in}} \cdot \beta.$$

Finally, we center the reading weights because the preceding layer norm projects out the all-ones vector. Thus, we center the weights $\mathbf{W}_{\text{eff}}$ as follows:

$$\mathbf{W}'_{\text{eff}}(i, :) = \mathbf{W}_{\text{eff}}(i, :) - \overline{\mathbf{W}}_{\text{eff}}(i, :).$$

Writing Weight Centering. Every time the model interacts with the residual stream, it applies a LayerNorm first. Therefore, the components of $\mathbf{W}_{\text{out}}$ and β_{out} that lie along the all-ones direction of

the residual stream have no effect on the model’s calculations. Consequently, we mean-center $\mathbf{W}_{\text{out}}$ and β_{out} by subtracting the means of the columns of $\mathbf{W}_{\text{out}}$:

$$\mathbf{W}'_{\text{out}}(:, i) = \mathbf{W}_{\text{out}}(:, i) - \overline{\mathbf{W}}_{\text{out}}(:, i).$$

Unembedding Centering. Since softmax is translation invariant, we also center $\mathbf{W}_U$:

$$\mathbf{W}'_U(:, i) = \mathbf{W}_U(:, i) - \overline{\mathbf{W}}_U(:, i).$$

D Experimental details

We carry out all our experiments using data from the C4 Corpus [63], which is publicly available through a ODC-BY license.⁹ We use the Pile [20], which is no longer distributed by its creators,¹⁰ solely to compute token frequency statistics, as Pythia was trained on it and we required the original training data for the model. The experiments to measure the total and direct effects, both for entropy and token frequency neurons, were carried out on 256-token input sequences. We used 100 sequences for GPT-2 and Pythia 410M, 50 for Phi-2 and GPT-2 Medium, and 30 for LLaMA2 7B and Gemma 2B. The scatter plot in Figure 3a was obtained on 10000 tokens. The results on the induction case study presented in §6.1 were obtained on 500 input sequences, the shaded area around each line represents the standard error. For the sake of visualization clarity, the box plots in Figure 2 and Figure 3 do not include outliers: the whiskers represent the first and third quartiles. For naturally occurring induction, we selected 1000 sequences from C4 that contained a repeated n-gram (n=6) and met our filtering criteria (each n-gram must be unique and contain no duplicate tokens). The BOS and mean attention output ablation experiments were each carried out on one hundred 100-token input sequences. Baselines for BOS ablation consisted of 6 single-head ablations, 3 two-head ablations, and 3 three-head ablations. Mean attention to BOS (used for baseline selection) was calculated using 1000 sequences from C4. To calculate prefix matching scores for attention heads, we use 1024 sequences of duplicated random tokens.

Computing resources. The static analyses of the model weights were conducted on a MacBook Pro with 32GB of memory. The experiments carried out to quantify the neurons’ effects were carried out on a single 80GB Nvidia A100. The longest experimental run took approximately 12 hours. Our experiments were carried out using PyTorch [59] and the TransformersLens library [53]. We performed our data analysis using NumPy [30] and Pandas [77]. Our figures were made using Plotly [37]. The paper’s bibliography was curated using Ryanize-bib [81].

E Unembedding null space & dropout

We hypothesize that the presence of a null space in the model’s unembedding matrix $\mathbf{W}_U$ is significantly affected by the application of dropout [66] during training. Dropout represents an incentive for the model to encode information redundantly, aligning it with a specific basis in the residual stream. We expect this redundancy to be reflected in some linear dependencies among the rows of $\mathbf{W}_U$, effectively creating a null space. To empirically validate this hypothesis, we compare the singular values of $\mathbf{W}_U$ between two versions of Pythia 160M [6] – identical except in their application of dropout to the residual stream (Figure 6).

Our observations reveal that the singular values of the model trained with dropout are generally lower, suggesting lower stability and greater linear dependence among the rows of $\mathbf{W}_U$. More notably, the difference in singular values between the two model versions is




Figure 6: Effect of Dropout on $\mathbf{W}_U$. Comparison of singular values for the unembedding matrix between two versions of Pythia 160M—one trained with dropout (red) and one without dropout (blue).

⁹<https://opendatacommons.org/licenses/by/1-0/>

¹⁰<https://the-eye.eu/public/AI/pile/readme.txt>




Figure 7: Entropy Neurons Across Different Models. Relationship between the fraction ρ of neuron norm in the $\mathbf{W}_U$ null space and the LayerNorm-mediated effect on model output for neurons in the last layer of various models. (a) Change in loss after ablation of entropy neuron 31.3228 in LLaMA2 7B. Ablation decreases entropy and increases loss on incorrect predictions.

larger for the ~ 10 smallest values, indicating the existence of a subspace in the dropout-trained model that has a substantially reduced impact on the logits. Such findings suggest that entropy neurons are more prevalent in models trained with dropout, as the subspace they leverage within the residual stream is more marked. However, this does not rule out the existence of entropy neurons in models trained without dropout.

F Entropy neurons in different models

In Figure 7, we report the results for the final-layer neurons of GPT-2 Medium, Pythia 410M and 1B, Phi-2, and Gemma 2B, complementing our earlier findings from GPT-2 and LLaMA2 7B discussed in §3. ρ is computed by considering the the bottom 1% $\mathbf{W}_U$ singular vectors as null space. We observe a consistent pattern: there exists a subset of neurons that predominantly write to the $\mathbf{W}_U$ null space. However, the extent to which entropy neurons appear in models appears to vary: we observe that entropy neurons in Pythia and Gemma exhibit a weaker effect compared to those in other models. In particular, the fraction of their total effect explained by the LayerNorm scale is around 30-40%, which is notably lower than the approximately 80% observed in GPT-2 Small, LLaMA2 7B, and Phi-2. This suggests that while entropy neurons are a common feature, their strength and the extent of their influence can vary significantly between different models.

F.1 Additional example of neuron activity

As for neuron 11.2378 in GPT-2 Small in §5, we analyze the strongest entropy neuron in LLaMA2 7B (31.3228). We perform a similar analysis as in §5, where we clip the activation value of the neuron back to its mean when it exceeds it, and we measure the changes in loss and entropy after the ablation (Figure 7a). The ablation of this neuron leads to a significant decrease in entropy, which results in a loss increase on incorrect predictions and a loss decrease on correct predictions.

G Token frequency neurons across different models

In Figure 8, we report the results comparing the token frequency-mediated effect against the average absolute change in the KL divergence between P_{model} and P_{freq} for neurons at the final layer of GPT-2




Figure 8: **Token Frequency Neurons Across Different Models.** Relationship between the token frequency-mediated effect and the average absolute change in the KL divergence from P_{freq} for final-layer neurons in (a) GPT-2 Small and (b) Pythia 1B. Neurons are categorized as normal, entropy, or token frequency neurons.




Figure 9: **GPT2-Small Entropy Neurons on Natural Induction:** (a) Change in activation values. (b) Change in entropy upon clipped mean ablation. (c) 11.2389 change in post-ablation loss. Color indicates reciprocal rank of the correct token prediction (1/RC). We use A0-A5 to refer to the positions of the repeated n-gram; B and X indicates the token position that follows the first and second occurrences of the n-gram

Small and Pythia 1B.¹¹ We also highlight the entropy neurons with the highest LayerNorm-mediated effect, represented by the points with the largest y-axis values in Figure 7.

In both models, we observe a set of neurons with substantial token frequency-mediated effect (~ 0.6). In Pythia 1B, the ablation of these neurons leads to a significant variation in the KL divergence from P_{freq} , comparable to the variations caused by entropy neurons. On the other hand, in GPT-2 Small, the effect of the entropy neurons on $D_{\text{KL}}(P_{\text{freq}}||P_{\text{model}})$ is significantly larger, suggesting that these neurons might play a more important role on the regulation of the model’s output compared to token frequency neurons.

Furthermore, in both models, we observe neurons with significant negative $1 - \text{DE}_{\text{freq}}/\text{TE}$ values, indicating that their effect increases when the impact of the token frequency direction is suppressed. This suggests that the effect through the token frequency direction actually counteracts their effect that bypasses the frequency direction.

H Additional Results on Induction

H.1 Naturally occurring induction

To verify that these neurons behave similarly outside of our artificial setting, we perform the same experiments on examples of induction in natural language data. We select 1000 sequences from the C4 corpus that contain a repeated n-gram of length $n=6$. In order to remove redundant and trivial examples, we also filter sequences such that the n-gram is unique and no individual n-grams contain duplicate tokens. We perform the same clipped mean ablation procedure as the synthetic setting, mean-ablating to the neuron’s activation value on random text.

¹¹We do not have access to the GPT-2 training corpus. However, since the model’s training data included a substantial amount of internet-scraped text, we use a randomly sampled 500M-token subset of the OpenWebText corpus [25] as a proxy for the original distribution.




Figure 10: **Token Frequency Neurons on Induction.** (a) Activations, entropy, and loss across duplicated 200-token input sequences. (b) The effect of clip mean-ablation of specific token frequency neurons. (c) Scatter plot of loss changes per token post-ablation of neuron 23.417, colored by reciprocal rank (RR). Ablation tends to increase loss for low initial loss tokens and decrease it for high initial loss tokens.

Figure 9 shows a similar trend to the synthetic setting shown in Figure 5a, though the mean change in activations (and thus change in entropy) is less significant. The most impactful neuron remains 11.2378, the ‘hedging’ neuron that we believe counter-balances overconfident predictions. Consistent with this hypothesis, when the neuron’s effect is removed via (clip) mean ablation, the model’s predictions become more confident: entropy decreases by up to 20% of its original value.

The other entropy-increasing neuron shown in Figure 9b is 11.2870. Unlike the other 4 entropy neurons, it fires strongly only on the first few tokens of induction and then decreases in activation (Figure 12). In both synthetic and natural induction settings, 11.2870 has an entropy-increasing effect (Figure 5b and Figure 9b), suggesting that it increases uncertainty at the start of induction where the model should be less certain that copying previous tokens will lead to the correct next-token prediction. (As an example, for the phrase ‘Mr. Smith ... Mr.’, the model should not be overly confident in predicting that the next token will be ‘Smith’, since a different name could follow instead.) Since we choose a relatively short prefix length for our natural induction setting ($n=6$), the neuron is active across our entire repeated sequence and increases entropy.

H.2 Token frequency neurons

To analyze the activity of token frequency neurons in the context of induction, we repeat the analysis carried out for entropy neurons. Using 100 input sequences created by concatenating two repetitions of a 100-token sequence, we study the activation values (averaged across token positions) for six token frequency neurons identified in §4 within the last-layer MLP in Pythia 410M. We observe that, in this setting, four neurons (417, 3412, 2017, and 87) display a substantial increase in activation value (Figure 10a). Upon clipped mean-ablation, neurons 417, 3412, and 87 lead to a significant increase in entropy compared to other neurons and a baseline of 10 randomly selected neurons (Figure 10b).

Focusing more closely on neuron 417, we find that reversing its state change leads to a somewhat positive delta in loss for tokens with initially low loss, and larger, negative changes in loss for tokens with high initial loss (Figure 10c). This trend, which is inverse to that observed with entropy neuron 2378 in GPT-2, demonstrates how different mechanisms of entropy modulation can operate at different levels to balance the model’s output distribution.

H.3 Change in activation of final-layer neurons

Figure 11 shows that entropy neurons exhibit among the largest changes in activation for final-layer neurons in GPT2 Small and LLaMa2 7B, suggesting that they are performing an important function in induction scenarios. When comparing synthetic and natural induction, we note that the set of most-changing entropy neurons is different (but overlapping). This is explained by the fact that most entropy neurons increase in activation over the course of induction (e.g. 11.2378), whereas others, such as 11.2870 only fire strongly at the start of induction. As discussed in Appendix H.1, our natural induction setting highlights start-of-induction neurons, since it consists of short repeated n -grams ($n=6$). By contrast, our synthetic setting consists of long 100-token repeated sequences and so favors neurons that continuously increase in activation.




Figure 11: **Neuron activation differences for synthetic and natural induction:** (left) GPT2-Small final-layer neurons (right) LLaMa2 7B final-layer neurons. Among final-layer neurons, entropy neurons exhibit the largest activation change for both synthetic and natural induction settings. Entropy neurons with especially high fraction of neuron norm in the null space (ρ) and large change in activation are labelled.

I Induction head–entropy neuron interaction

In this section, we study the interactions between induction heads and entropy neurons in more detail. In Appendix I.1 we use BOS ablations to verify that induction heads have a significant causal effect for 5/6 of our entropy neurons. Next, in Appendix I.2 we provide preliminary evidence that induction heads produce an internal signal that indicates when induction has occurred, and that our entropy neurons respond to this signal. Finally, in Appendix I.3 we check the direct composition between entropy neurons and induction heads, showing that induction heads do not have especially strong direct effect, and identifying a set of components that may mediate the ‘induction-has-occurred’ signal.

I.1 BOS Ablations for GPT2 Small neurons

We present the full results of BOS ablations (§6.2) for all six of the entropy neurons under investigation in GPT2 Small (Figure 12). Neurons 11.584, 11.2123, and 11.2910 display a similar trend to 11.2378: intervening on the top induction head, L5H1, leads to a substantial decrease in activation, suggesting it has an important causal effect. By comparison, ablating baseline heads in layers 5 and 6 have little impact on activation.

We note that for 11.2870, our start-of-induction neuron (Appendix H.1), BOS ablation of L5H1 has a strong causal effect, but that the direction of this effect changes: on the first few tokens of the repeated sequence, the head boosts activation (activation decreases upon BOS ablation) but it suppresses activation across the remainder of the sequence (activation increases upon ablation). Prior work has found that attention head L5H1 is specialized for long-prefix induction (induction over long sequences) [44]. Thus, we believe 11.2870 is a concrete example of where such specialization is useful: the output of L5H1 enables 11.2870 to fire strongly (and increase entropy) only at the start of induction.

I.2 Mean attention output ablations

Because entropy neurons respond strongly to induction and are causally influenced by induction heads, we hypothesize that induction heads produce a generic ‘induction has occurred’ signal that




Figure 12: **GPT2-Small Entropy Neurons after BOS Ablations:** Ablation of induction heads leads to substantial decrease in activation for most entropy neurons.




Figure 13: **GPT2 Small Entropy Neurons after Mean Attention Ablations:** Replacing induction head outputs with their mean induction output leads to similar activations.

indicates to later components that the model is in an induction context. To test this hypothesis, in our synthetic induction setting, we perform mean attention output ablations: we replace each induction head’s contribution to the residual stream with its mean contribution averaged over the second half of the sequence (i.e. when induction is occurring).¹² In this way, we hope to retain the generic “induction has occurred” signal, while removing token-specific behavior, such as directly boosting the logits of an attended-to token.

Figure 13 shows that mean attention ablations are able to produce activations that are close to their original values. For example, for 11.2378, mean attention ablation of the top induction head results in small activation decrease from ~ 9 to ~ 8 . For comparison, recall that BOS ablation of this same head causes mean activation to fall to roughly 4. This suggests that a substantial portion of the

¹²In practice, this is implemented using attention $\mathbf{Z}$ vectors (‘attention output pre-linear’) rather than attention output. However, the two are equivalent for our purposes since $\mathbf{Z}$ vectors are converted to attention outputs by a linear map [53]




Figure 14: **GPT2 Small Direct Effect of attention heads and mlps on induction: DFA difference** between induction and non-induction sequences. Attention heads are to the left of the dotted line. Induction score refers to the prefix matching score.

entropy neurons' activation is derived from this generic induction signal, rather than token-specific information.

We note that though mean attention ablation is sufficient to raise the activations of our entropy neurons close to their original activations on the second half of the sequence, they do not have much impact over the first half of the sequence. (We would expect activation to increase, since we are adding an 'induction is happening' signal.) We speculate that this is because there are other important components that we have not factored into our analysis. For example, the signal from induction heads to entropy neurons is likely mediated by later-layer MLPs/Attention Heads. These could perform an AND operation over our induction heads' signals and duplicate token heads, such that both signals are necessary in order to produce the trigger for our entropy neurons. We leave a more thorough investigation of this potential circuit to future work.

I.3 Direct effect of components

In this section, we study whether the output of induction heads directly contributes to the activation of our entropy neurons, or whether their effect is mediated by other components. To do this, we produce a corrupted version of our induction sequences by replacing the first 100 tokens of each sequence with another random 100-token sequence sampled from C4. Then, for both our original and corrupted sequences, we use direct feature attribution (DFA) [43] to calculate the direct contribution of each preceding component in the transformer to our neuron's activation across the last 100 tokens. This is done by taking the dot product of the output of each component (attention heads and MLPs) with the neuron's $\mathbf{W}_{in}$ vector, after applying LayerNorm.

Figure 14 shows the difference in DFA between the original (induction) sequence and the corrupted (non-induction) sequence, indicating which components were responsible for the change in activation. For entropy neuron 11.2870 induction heads have a strong direct contribution to the neuron's activation. However, this is not the case for our other 5 entropy neurons. In particular, L11H8 as well as MLPs in layers 8-10 have a strong direct effect, suggesting that the 'induction has occurred' signal effect may be mediated by these other components. Further evidence for this is suggested by the observation that the direct effect of induction heads is often opposite to the expected change in activation. That is to say, for entropy neurons that increase in activation upon induction, the direct effect of the induction heads is often negative (i.e. suppressing activation), and becomes even more negative upon induction. Since our BOS ablation experiments show that turning off these heads leads to a *decrease* in activation, this suggests the presence of intermediate components that convert the negative direct effect signal from induction heads into a positive one.

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: §3 and §4.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: Appendix A

Guidelines:

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

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We do not provide any theoretical result.

Guidelines:

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

4. Experimental Result Reproducibility

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

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

Justification: §3, §4, and Appendix D

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: We uploaded our code through the submission platform and we provide experimental details in Appendix D.

Guidelines:

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

6. Experimental Setting/Details

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

Answer: [Yes]

Justification: §3, §4, and Appendix D

Guidelines:

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

7. Experiment Statistical Significance

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

Answer: [Yes]

Justification: In §3 and §4 we include box plots, which provide comprehensive information about the distribution of the effects measured. For the line plots in section §6, we indicate the standard error as a shaded area around each line.

Guidelines:

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

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

8. Experiments Compute Resources

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

Answer: [Yes]

Justification: Appendix D

Guidelines:

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

9. Code Of Ethics

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

Answer: [Yes]

Justification: Appendix B

Guidelines:

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

10. Broader Impacts

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

Answer: [Yes]

Justification: Appendix B

Guidelines:

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

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

11. Safeguards

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

Answer: [NA]

Justification: We do not release any data or model with a high risk of misuse.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: Appendix D.

Guidelines:

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

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

13. **New Assets**

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

Answer: [NA]

Justification: We do not release new assets.

Guidelines:

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

14. **Crowdsourcing and Research with Human Subjects**

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

Answer: [NA]

Justification: We do not involve crowdsourcing nor research with human subjects.

Guidelines:

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

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

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

Answer: [NA]

Justification: We do not involve crowdsourcing nor research with human subjects.

Guidelines:

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