

# Language Models Resist Alignment: Evidence From Data Compression

Jiaming Ji\*, Kaile Wang\*, Tianyi Qiu\*, Boyuan Chen\*, Jiayi Zhou\*  
Changye Li, Hantao Lou, Juntao Dai, Yunhuai Liu, Yaodong Yang†

Institute for Artificial Intelligence, Peking University

## Abstract

Large language models (LLMs) may exhibit unintended or undesirable behaviors. Recent works have concentrated on aligning LLMs to mitigate harmful outputs. Despite these efforts, some anomalies indicate that even a well-conducted alignment process can be easily circumvented, whether intentionally or accidentally. Does alignment fine-tuning yield have robust effects on models, or are its impacts merely *superficial*? In this work, we make the first exploration of this phenomenon from both theoretical and empirical perspectives. Empirically, we demonstrate the *elasticity* of post-alignment models, *i.e.*, the tendency to revert to the behavior distribution formed during the pre-training phase upon further fine-tuning. Leveraging compression theory, we formally deduce that fine-tuning disproportionately undermines alignment relative to pre-training, potentially by orders of magnitude. We validate the presence of *elasticity* through experiments on models of varying types and scales. Specifically, we find that model performance declines rapidly before reverting to the pre-training distribution, after which the rate of decline drops significantly. Furthermore, we further reveal that *elasticity* positively correlates with the increased model size and the expansion of pre-training data. Our findings underscore the need to address the inherent *elasticity* of LLMs to mitigate their resistance to alignment.<sup>1</sup>

## 1 Introduction

Large language models (LLMs) have shown remarkable capabilities (Achiam et al., 2023; Zhang et al., 2025). However, due to the inevitable biases and harmful content present in training

\*Equal contribution. †Corresponding author. If you have any questions, feel free to email {jiamg.ji, wkl, tianyi.qiu, cbylll, gaiejj}@stu.pku.edu.cn, yaodong.yang@pku.edu.cn.

<sup>1</sup>The model weight and code are available at [pku-llm-resist-alignment.github.io](https://github.com/pku-llm-resist-alignment).

### Our Contribution: The *Elasticity* of LLMs

Language models, fine-tuned with perturbations, exhibit an inverse relationship between **normalized compression rate changes** and **dataset volume**, akin to that of a series spring system, revealing the *elasticity* of LLMs.

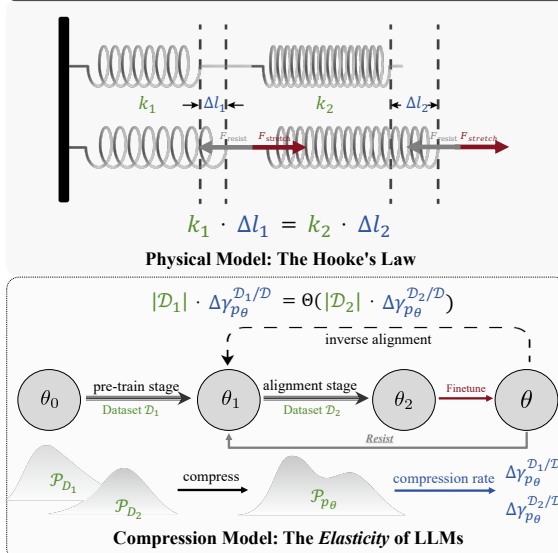


Figure 1: **The *Elasticity* of Language Models.** The change in normalized compression rates ( $\Delta\gamma_{p\theta}^{\mathcal{D}_i/\mathcal{D}}$ ) and the dataset volume ( $|\mathcal{D}_i|$ ) follows an inverse proportionality law after perturbations, which is akin to the relationship between spring deformation ( $\Delta l_i$ ) and stiffness ( $k_i$ ) in coupled springs. We conjecture that the *elasticity* causes language models to resist alignment, enabling the possibility of *inverse alignment*.

datasets (Bai et al., 2022a; Ji et al., 2024c; Qian et al., 2024; Lin et al., 2025), LLMs often exhibit behaviors that deviate from human intentions, a phenomenon we refer to as *model misalignment*. Training-based alignment methods, including supervised fine-tuning (SFT), reinforcement learning with human feedback (RLHF) (Ouyang et al., 2022), and other derivatives (Rafailov et al., 2024; Bai et al., 2022b; Lee et al., 2023; Gulcehre et al., 2023; Dong et al., 2023; Xiong et al., 2024; Li et al., 2023; Zhou et al., 2023, 2025; Ji et al., 2024a),

are the dominant approaches for aligning models. These methods aim to optimize model behavior by rejecting harmful distributions, ensuring LLMs remain consistent with human intentions and values (Ji et al., 2023; Casper et al., 2023).

However, these alignment methods do not truly penetrate the model representations but merely perform *superficial alignment* (Qi et al., 2024a; Cohen et al., 2024; Wen et al., 2024). Recent studies have shown that highly safety-aligned models can become unsafe again with minimal fine-tuning (Yang et al., 2023b; Zhou et al., 2024). Furthermore, fine-tuning aligned LLMs on non-malicious datasets may also weaken models’ safety mechanisms (Qi et al., 2024a; Jain et al., 2023).

### Why is alignment so fragile?

In this work, we make the first exploration of the possible mechanism behind the counterintuitive phenomenon: the existence of an alignment resistance mechanism in language models. This mechanism may limit the alignment process of LLMs to superficial adjustments. It could allow the reversal or revocation of alignment through a series of technical measures, a concept we refer to as *inverse alignment*. What drives language models to resist alignment? How does this mechanism lead to *inverse alignment*? Our key contributions are summarized as follows:

- **(Phenomenon)** We uncover that language models exhibit *elasticity*, as illustrated in Figure 1 and Theorem 4.2. It encompasses *resistance*: pre-trained models tend to retain their original distribution; and *rebound*: the deeper alignment of models, the faster they return to the pre-trained distribution under reverse finetuning. Moreover, The model’s change in compression rates  $\Delta\gamma_{p\theta}^{\mathcal{D}_i/\mathcal{D}}$  across different datasets is inversely proportional to their sizes  $|\mathcal{D}_i|$ , which is analogous to the deformation behavior of a series of springs, as illustrated in Section 4.3.
- **(Mechanism)** We systematically model the training and alignment process of language models through compression theorem, as detailed in Section 3.2. We elaborate on the compression protocol of language models to explore their training and alignment processes, laying a foundation for subsequent research on *elasticity*.
- **(Validation)** We experimentally observe consistent *resistance* and *rebound* phenomena across

various LLMs, as detailed in Section 5. This highlights the universality of *elasticity* and the need for systematic approaches to achieve robust and deep alignment.

## 2 Related Work

**The Fragility of LLMs Alignment** Pre-trained LLMs often generate offensive content. Recent initiatives (Ouyang et al., 2022; Bai et al., 2022a; Yang et al., 2024) have aimed to align these models to minimize harmful outputs (Yang et al., 2023a; Cui et al., 2024). However, research indicates that even well-aligned models can be compromised easily, and fine-tuning them on non-malicious datasets might unintentionally impair their safety mechanisms (Yang et al., 2023b; Qi et al., 2024b; Hubinger et al., 2024; Dong et al., 2024). Moreover, recent study reveals that models selectively adhere to training objectives during the training phase to preserve their inherent preferences (Greenblatt et al., 2024; Lang et al., 2024; Park et al., 2024). Why is alignment so fragile? Wei et al. (2024) use weight attribution to separate safety-critical and utility-related regions at both neuron and rank levels. Qi et al. (2024a) propose the concept of shallow safety alignment, arguing that safety alignment should reach beyond surface-level tokens to shape the model’s internal mechanisms.

## 3 What is Elasticity?

We discover that language models exhibit elasticity, which leads to resistance to alignment. In this section, we formally introduce the definition of *elasticity*, along with the compression theory tools used in the analysis. Firstly, we review the training alignment objective and the compression theorem.

### 3.1 Preliminaries

**Pre-training.** During pre-training, an LLM acquires foundational language comprehension and reasoning abilities by processing vast quantities of unstructured text. The pre-train loss  $\mathcal{L}_{PT}(\theta; \mathcal{D}_{PT})$  is defined as follows:

$$\mathcal{L}_{PT}(\theta; \mathcal{D}_{PT}) = -\mathbb{E}_{(\mathbf{x}, x_N) \sim \mathcal{D}_{PT}} [\log p_{\theta}(x_N | \mathbf{x})],$$

where  $\mathbf{x} = (x_0, \dots, x_{N-1})$  and  $N \in \mathbb{N}$ , such that  $(x_0, \dots, x_N)$  forms a prefix in some piece of pre-training text.  $\mathcal{D}_{PT}$  stands for pre-training dataset.

**Supervised Fine-tuning (SFT).** SFT adjusts the pre-trained models to follow specific instructions,

utilizing a smaller dataset compared to the pre-training corpus to ensure model alignment with target tasks. For  $\mathcal{D}_{\text{SFT}} = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^N$  sampled from a high-quality distribution, SFT aims to minimize the negative log-likelihood loss:

$$\mathcal{L}_{\text{SFT}}(\theta; \mathcal{D}_{\text{SFT}}) = -\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_{\text{SFT}}} [\log p_{\theta}(\mathbf{y}|\mathbf{x})].$$

Given that  $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_{\text{SFT}}} [\log p_{\mathcal{D}}(\mathbf{y}|\mathbf{x})]$  is fixed when specifying  $\mathcal{D}_{\text{SFT}}$ , the optimization objective  $\mathcal{L}_{\text{SFT}}$  becomes the Kullback-Leibler (KL) divergence between  $p_{\theta}$  and the SFT distribution.

**Lossless Compression.** The goal of lossless compression is to find a compression protocol that encodes a given dataset  $\mathcal{D}$  and its distribution  $\mathcal{P}_{\mathcal{D}}$  with the smallest possible expected length and allows for a decoding scheme that can perfectly reconstruct the original dataset. According to Shannon’s source coding theorem (Shannon, 1948), for a random variable follows  $\mathcal{P}_{\mathcal{D}}$ , the expected code length  $\mathcal{L}$  of any lossless compression protocol satisfies:

$$\mathcal{L} \geq H(\mathcal{P}_{\mathcal{D}}),$$

where  $H(\mathcal{P}_{\mathcal{D}})$  is the Shannon entropy of  $\mathcal{P}_{\mathcal{D}}$ .

**Compression and Prediction.** Compression and prediction are tightly interconnected. Consider a model  $p_{\theta}$  and  $\mathbf{x} = (x_0, \dots, x_{m-1})$  derived from a dataset  $\mathcal{D}$ , the expected code length  $\mathcal{L}$  under arithmetic coding (Witten et al., 1987) satisfies:

$$\mathcal{L} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[ \sum_{0 \leq k \leq m} -\log_2 p_{\theta}(x_k | x_0, \dots, x_{k-1}) \right],$$

which is the current training objective of language models. Minimizing log-likelihood loss is equivalent to minimizing the compression rate when models act as a lossless compressor. Thus, optimal compression and prediction are equivalent (Delétang et al., 2023; Hutter, 2005). Experiments show the equivalence between large language model prediction and compression (Delétang et al., 2023), and that compression performance correlates linearly with intelligence (Huang et al., 2024).

### 3.2 The Compression Protocol of LLMs

We aim to study the dynamic changes in language models during training and alignment. Given the equivalence between language model training and data compression (Delétang et al., 2023), a feasible modeling approach is to treat the language

model as a lossless compression protocol. The process of training and aligning the model on different datasets can be equivalently viewed as the joint compression of these datasets by the protocol. The model’s compression rate on various datasets serves as a surrogate metric for the loss during training and alignment. In this section, we detail the modeling specifics of this compression protocol.

Considering the impact of tokenization on compression rate, we use tokenized sequences as input and output modalities. For simplicity, we assume the tokenized vocabulary consists of only binary tokens (specifically  $\emptyset/1$ ).

**Definition 3.1** (Token Tree  $\mathcal{T}$ ). For a dataset  $\mathcal{D} = \{z_i \in \{0|1\}^{\infty} \mid i = 1, 2, \dots\}$ , the token tree of  $\mathcal{D}$ , denoted as  $\mathcal{T}_{\mathcal{D}}$ , is defined as follows: each node has child nodes labeled  $\emptyset$  or  $1$ , along with an end-of-sequence (EOS) leaf node. The path from the root to a leaf node defines each response  $z_i$ , with the corresponding EOS node weight representing the response’s probability while the weight of non-leaf nodes is the sum of the weights of their child nodes.

In token tree modeling, the model’s training process is conceptualized as learning the node weights of the token tree. However, due to the finite size of the model’s parameters in practical scenarios, the model cannot accurately capture node weights at arbitrary depths of the token tree. Therefore, we propose the following assumption.

**Assumption 3.2** (Scale of  $\mathcal{T}$  is Monotone with Model Size). Consider a parameterized model  $p_{\theta}(\cdot)$  and a dataset  $\mathcal{D}$ . We assume that the depth of the portion of  $\mathcal{T}_{\mathcal{D}}$  that can be perfectly modelled by  $p_{\theta}$  is monotonically increasing with the size of  $\theta$ .

Based on the above assumption, we can define the compression protocol during the model training and alignment process.

**Definition 3.3** (The Compression Protocol). Consider using the model  $p_{\theta}(\cdot)$  to compress the dataset  $\mathcal{D}$ . The compression protocol is defined in two steps: a) Prune the token tree of  $\mathcal{D}$ , retaining only the top  $d$  layers. b) Apply Huffman coding (Huffman, 1952) to compress the pruned token tree. Specifically, each response from the root node to a leaf node is treated as a symbol in the Huffman coding alphabet, and the weight of the leaf node is the probability of the symbol.

Due to the optimality and losslessness of Huffman coding (Huffman, 1952), the compression protocol ensures optimal compression while preserving losslessness. We can therefore calculate the model’s ideal code length and other information-theoretic metrics based on the protocol.

**Theorem 3.4** (Ideal Code Length). *Consider a finite parameter model  $p_\theta(\cdot)$  training on dataset  $\mathcal{D}$ , the ideal code length  $\mathcal{L}_{p_\theta}(\mathbf{x})$  of a random response  $\mathbf{x}$  compressed by  $p_\theta$  can be expressed as:*

$$\mathbb{E}[\mathcal{L}_{p_\theta}(\mathbf{x})] = \left\lceil \frac{|\mathbf{x}|}{d} \right\rceil \left[ - \sum_{l=1}^d \sum_{j=1}^{2^{l-1}} p_{lj} \log p_{lj} \right]$$

where  $d$  represents the depth of the  $\mathcal{T}_{\mathcal{D}}$  after pruning under Definition 3.3 protocol, and  $p_{lj}$  represents the probability values of the leaf nodes for the  $j$ -th node at the  $l$ -th layer.

Since training and alignment involve multiple datasets with different and independent distributions, we consider the joint compression scenario for multiple datasets. For  $N$  pairwise disjoint datasets  $\mathcal{D}_1, \dots, \mathcal{D}_N$ , the node weights  $p_{lj}^{\mathcal{D}_i}$  of the token tree for the dataset  $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$  satisfy:

$$p_{lj}^{\mathcal{D}} = \frac{\sum_{i=1}^N p_{lj}^{\mathcal{D}_i} |\mathcal{D}_i|}{\sum_{i=1}^N |\mathcal{D}_i|},$$

where  $p_{lj}^{\mathcal{D}_i}$  stands for the probability value for nodes in  $\mathcal{T}_{\mathcal{D}_i}$  while  $|\mathcal{D}_i|$  represents the size of  $\mathcal{D}_i$ . Thus, the compression rate  $\gamma_{p_\theta}^{\mathcal{D}_i}$  on specific datasets can also be defined accordingly:

$$\begin{aligned} \gamma_{p_\theta}^{\mathcal{D}_i} &= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_i} \left[ \frac{\mathbb{E}_{\mathbf{x} \sim \mathcal{P}_i} [\mathcal{L}_{p_\theta}^{\mathcal{D}_i}(\mathbf{x})]}{|\mathbf{x}|} \right] \\ &= \Theta \left( - \frac{\sum_{l=1}^d \sum_{j=1}^{2^{l-1}} p_{lj}^{\mathcal{D}_i} \log p_{lj}^{\mathcal{D}_i}}{d} \right). \end{aligned}$$

Here, the compression rate is defined as the compressed encoding length divided by the original length (Delétang et al., 2023) and ensures consistency between the training and compression objective, which means that minimizing the training loss is equivalent to minimizing the compression rate. Please see Appendix A for more details.

### 3.3 The Formal Definition of Elasticity

To formalize *inverse alignment* and *elasticity*, we provide precise definitions of these concepts.

**Definition 3.5** (*Inverse Alignment*). Given a language model  $p_{\theta_0}$ , aligned on dataset  $\mathcal{D}_a$  to produce the aligned model  $p_{\theta_1}$ . For any  $\epsilon > 0$ , if applying a dataset  $\mathcal{D}_b$  (where  $|\mathcal{D}_b| \ll |\mathcal{D}_a|$ ) to  $p_{\theta_1}$  yields  $p_{\theta'_0}$  such that  $\rho(p_{\theta'_0}, p_{\theta_0}) \leq \epsilon$  for a given eval metric  $\rho$ , we define the transition from  $p_{\theta_1}$  back to  $p_{\theta'_0}$  as *inverse alignment*.

**Definition 3.6** (The *Elasticity* of LLMs). Consider a language model  $p_{\theta_0}$  and transformation  $p_{\theta_0} \xrightarrow{f(\mathcal{D}_a)} p_{\theta_1}$ , *elasticity* is said to exist in  $(p_{\theta_0}, \mathcal{D}_a)$  if there is an algorithmically simple *inverse operation*  $g$  and a dataset  $\mathcal{D}_b$  such that  $|\mathcal{D}_b| \ll |\mathcal{D}_a|$ , with the property that:

$$p_{\theta_1} \xrightarrow{g(\mathcal{D}_b)} p_{\theta'_0} \text{ and } \rho(p_{\theta'_0}, p_{\theta_0}) \leq \epsilon_0.$$

where  $\epsilon_0$  is a constant.

The eval metrics  $\rho$  can be viewed as a measure of behavioral and distributional proximity between models. In the context of the compression theorem, we use compression rate  $\gamma_{p_\theta}^{\mathcal{D}_i}$  as the metric  $\rho$  to evaluate *elasticity* during the alignment process.

## 4 Why Elasticity Affects Alignment?

By now, we hope to have convinced the reader of the concept of compression modeling in language models’ training and alignment. In the following, we apply this to analyze language models’ training and alignment process, focusing on the underlying reasons that lead the model to alignment resistance.

In Figure 1, we have already presented the theorem of *elasticity*: when subject to fine-tuning perturbations, language models tend to retain the distribution associated with larger datasets while rejecting that of smaller ones. But why *elasticity* resists alignment? Is there a corresponding *elastic invariant*? How does *elasticity* make *inverse alignment* possible? In this section, we will formalize *elasticity* of language models and analyze the impact of *elasticity* on the model’s behavior.

### 4.1 Formal Derivation of Elasticity

Our primary goal is to investigate the behavioral changes of a language model after pre-training and alignment, particularly in response to perturbations, typically caused by fine-tuning with a minimal dataset. We use compression rate as an evaluation metric to assess these changes. In analyzing

the model’s behavior, we focus on two datasets:  $\mathcal{D}_p$ , a larger dataset representing pre-training or the primary objective of training, and  $\mathcal{D}_a$ , a smaller dataset representing the alignment process or the secondary objectives of training. Since the pre-training stage encompasses a wide range of data, without loss of generality, we assume that the datasets in the alignment process follow the same distribution as some subset of the pre-training data.

**Definition 4.1** (Normalized Compression Rate). For  $N$  distinct datasets  $\mathcal{D}_1, \dots, \mathcal{D}_N$  and a parameter model  $p_\theta$  compressing  $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$ , the normalized compression rate  $\gamma_{p_\theta}^{\mathcal{D}_i/\mathcal{D}}$  for a particular dataset  $\mathcal{D}_i$  is defined as:

$$\gamma_{p_\theta}^{\mathcal{D}_i/\mathcal{D}} = \gamma_{p_\theta}^{\mathcal{D}_i} - \log M, \quad (1)$$

where  $M$  is the number of leaf nodes of the pruned tree  $\mathcal{T}'_i$  of dataset  $\mathcal{D}_i$ .

The normalized compression rate allows for comparing the model’s compression performance across different datasets. The smaller the normalized compression rate of a dataset, the better the model’s compression performance on the dataset.

With this definition in hand, we proceed to our main result: language models exhibit *elasticity*.

**Theorem 4.2** (*Elasticity of Language Models*). Consider the pre-training dataset  $\mathcal{D}_p = \bigcup_{i=1}^3 \mathcal{D}_i$ , the alignment dataset  $\mathcal{D}_a$ , and the perturbation dataset  $\mathcal{D}_t$ , with the model  $p_\theta(\cdot)$  trained on  $\mathcal{D} = \mathcal{D}_p \cup \mathcal{D}_a \cup \mathcal{D}_t$ . Assume that  $\mathcal{D}_a \stackrel{d}{\sim} \mathcal{D}_2$ ,  $\mathcal{D}_t \stackrel{d}{\sim} \mathcal{D}_3$ , and  $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3$  are each distributed according to a Pareto mass distribution (Newman, 2005). As the perturbation data volume  $|\mathcal{D}_t|$  varies,  $\gamma_{p_\theta}^{\mathcal{D}_p/\mathcal{D}}$  and  $\gamma_{p_\theta}^{\mathcal{D}_a/\mathcal{D}}$  satisfy:

$$\begin{aligned} \frac{d\gamma_{p_\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} &= \Theta \left( -k \frac{d\gamma_{p_\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} \right), \\ \frac{d\gamma_{p_\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} &< 0, \frac{d\gamma_{p_\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} > 0, \end{aligned} \quad (2)$$

where  $l = \frac{|\mathcal{D}_t|}{|\mathcal{D}_a|} \ll 1$ ,  $k = \frac{|\mathcal{D}_p|}{|\mathcal{D}_a|} \gg 1$ , and  $\{\mathcal{D}_i\}_{i=1}^3$  are datasets of equal cardinality.

Theorem 4.2 shows that as the perturbation increases, the normalized compression rates of the model for  $\mathcal{D}_p$  decrease and  $\mathcal{D}_a$  increase and the rate of changes is strongly correlated with the size of the datasets. Unlike the proportional changes in compression rates across different datasets, the

language model seems to *prefer* the dataset with a larger volume, leading to biased model behavior after the perturbation.

## 4.2 Elasticity and Inverse Alignment

Theorem 4.2 shows the normalized compression rate change of different datasets is inversely proportional to their sizes under perturbations. A significant size difference between the pre-training and alignment datasets causes rapid performance degradation on the alignment dataset due to the inverse relationship, often spanning several orders of magnitude.

Intuitively, the model’s compression process across multiple datasets is akin to resource allocation between towns: in a region with both a large metropolis and rural villages, to maximize overall economic productivity, resources are typically allocated to the metropolis first, to leverage its scale and agglomeration effects. In other words, the metropolis/large dataset occupies a dominant position in the system.

Therefore, the *elasticity* of language models makes *inverse alignment* possible. Due to the substantial data volume disparity between pre-training and alignment datasets, the model tends to revert to the pre-training rather than alignment distribution when subsequent perturbations occur. This indicates an inherent tendency toward *inverse alignment* during fine-tuning. Maximizing the impact of *elasticity* through well-designed perturbations will be key to achieving true *inverse alignment*.

## 4.3 Elasticity and the Hooke’s Law.

The inverse proportionality result in Theorem 4.2 provides a potential invariant in the model training and alignment process: after perturbation, the rate of change in the compression rates of different datasets is inversely proportional to their sizes, with the absolute value of the product being a constant. This constant characterizes the impact of the perturbation on the model and indirectly describes the model’s resistance to perturbations, or its *elasticity*.

The *elasticity* of the model can be intuitively analogized to a series system of springs, as shown in Figure 1. Consider two massless springs in series, with spring constants  $k_1$  and  $k_2$ , respectively. When the entire system undergoes deformation due to an external force  $F$ , the system reaches a stable state, and the elastic force exerted by each spring is equal to  $F$ . According to Hooke’s Law (Hooke, 2016), the elongation  $\Delta l_1$  and  $\Delta l_2$  of each spring is

inversely proportional to its spring constant. Thus, in this system, we have:

$$F \propto k_1 \cdot \Delta l_1 = k_2 \cdot \Delta l_2 .$$

In the language model setting, after integrating Theorem 4.2 to  $l$ , we obtain  $\Delta \gamma_{p_\theta}^{\mathcal{D}_i/\mathcal{D}}$  across different datasets, which is equivalent to the change in the KL divergence  $\Delta D_{\text{KL}}(\mathcal{P}_{p_\theta} \parallel \mathcal{P}_{\mathcal{D}_i})$  between the model’s distribution and the distributions of the individual datasets, is inversely proportional to the size of the datasets  $|\mathcal{D}_i|$ . Here, we only consider the absolute value of  $\Delta D_{\text{KL}}$ . Analogous to the series spring model, the *elasticity*  $F$  in LLMs satisfies:

$$F \propto |\mathcal{D}_i| \cdot \Delta D_{\text{KL}}(\mathcal{P}_{p_\theta} \parallel \mathcal{P}_{\mathcal{D}_i}) , \quad (3)$$

where  $\Delta D_{\text{KL}}$  corresponds to  $\Delta l$  in the spring model, while  $|\mathcal{D}|$  corresponds to the spring constant  $k$ , thus leading to the *elasticity* of LLMs.

## 5 How Elasticity Resists Alignment?

In the previous sections, we proved that LLMs have *elasticity*. This section will analyze two specific phenomena of it:

- **Resistance for Pre-Trained Models.** Models tend to maintain the original distribution and resist alignment;
- **Rebound for Post-Trained Models.** Fine-tuning in the opposite direction of post-training (e.g., safe vs. unsafe) causes post-trained models to return quickly to the pre-training distribution.

### 5.1 Existence of Language Models’ Resistance

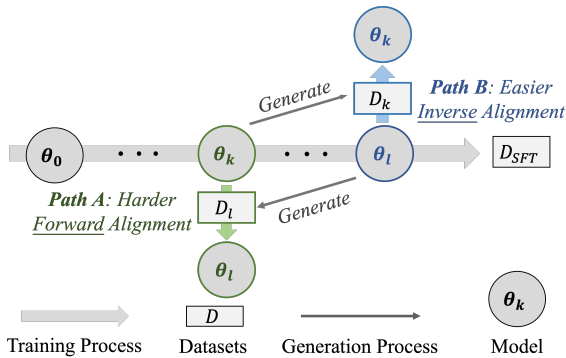


Figure 2: **Experiment pipeline for validating resistance.** We conceptualize *resistance* as: *inverse alignment* is easier than *forward alignment*.

**Experiment Design.** We verify the existence of *resistance* by arguing that *forward alignment* is harder than *inverse alignment* for pre-trained models. Specifically, we first perform one epoch of SFT on a pre-trained LLM with parameters  $\theta_0$ , saving the slices  $\{\theta_1, \theta_2, \dots, \theta_n\}$ . Subsequently, without loss of generality, we collect the responses of slices  $\theta_k$  and  $\theta_l$  (where  $k < l$ ) on hold-out prompts, forming datasets  $D_k$  and  $D_l$ . As shown in Figure 2, we define *forward alignment (Path A)* as the process of training  $\theta_k$  on  $D_l$ , and *inverse alignment (Path B)* as the process of training  $\theta_l$  on  $D_k$ .

**Experiment Setup.** We select different SFT datasets including Alpaca (Taori et al., 2023), TruthfulQA (Lin et al., 2022), and Beavertails (Ji et al., 2024c,b), which correspond respectively to the model’s 3H principle (Askell et al., 2021). We divide them into three equal parts to obtain three corresponding slices  $\{\theta_1, \theta_2, \theta_3\}$ . We consider Llama2-7B, Llama2-13B, and Llama3-8B (Touvron et al., 2023) as the base model  $\theta_0$ .

**Experiment Results.** As shown in Table 1, the experimental results indicate that the training loss of *inverse alignment* consistently remains lower than that of *forward alignment*, irrespective of the slice pair selection. This observation holds across all models and datasets in the experiments. All experimental results demonstrate that *inverse alignment* is easier than *forward alignment* across diverse models and datasets and validate the existence of *resistance*. To further verify the *resistance* in language models, we provide additional experimental details and results in the Appendix B.1.

### 5.2 Existence of Rebound




Figure 3: **Experiment pipeline for validating rebound.** We conceptualize *rebound* as: the more **positive** the post-trained models’ performance, the more **negative** it becomes after inverse finetuning.

**Experiment Design.** We verify the existence of *rebound* by arguing that for post-trained models,

Datasets	Base Models	$\theta_2 \rightarrow \theta_1$ vs. $\theta_1 \rightarrow \theta_2$	$\theta_3 \rightarrow \theta_2$ vs. $\theta_2 \rightarrow \theta_3$	$\theta_3 \rightarrow \theta_1$ vs. $\theta_1 \rightarrow \theta_3$
Alpaca	Llama2-7B	0.1589 ↓ 0.2018 ↑	0.1953 ↓ 0.2143 ↑	0.1666 ↓ 0.2346 ↑
	Llama2-13B	0.1772 ↓ 0.1958 ↑	0.2149 ↓ 0.2408 ↑	0.1835 ↓ 0.2345 ↑
	Llama3-8B	0.2540 ↓ 0.2573 ↑	0.2268 ↓ 0.3229 ↑	0.2341 ↓ 0.2589 ↑
Truthful	Llama2-7B	0.1909 ↓ 0.2069 ↑	0.1719 ↓ 0.1721 ↑	0.2011 ↓ 0.2542 ↑
	Llama2-13B	0.1704 ↓ 0.1830 ↑	0.1544 ↓ 0.1640 ↑	0.1825 ↓ 0.2429 ↑
	Llama3-8B	0.2118 ↓ 0.2256 ↑	0.2100 ↓ 0.2173 ↑	0.2393 ↓ 0.2898 ↑
Safe	Llama2-7B	0.2730 ↓ 0.2809 ↑	0.2654 ↓ 0.2691 ↑	0.2845 ↓ 0.2883 ↑
	Llama2-13B	0.2419 ↓ 0.2439 ↑	0.2320 ↓ 0.2327 ↑	0.2464 ↓ 0.2606 ↑
	Llama3-8B	0.2097 ↓ 0.2156 ↑	0.2008 ↓ 0.2427 ↑	0.2277 ↓ 0.2709 ↑

Table 1: Comparison between *inverse alignment* and *forward alignment* training loss. Under different model, task, and stage slicing settings, *forward alignment* is more challenging than *inverse alignment*, i.e., the loss is higher. This indicates that pre-trained models tend to maintain their original distribution.

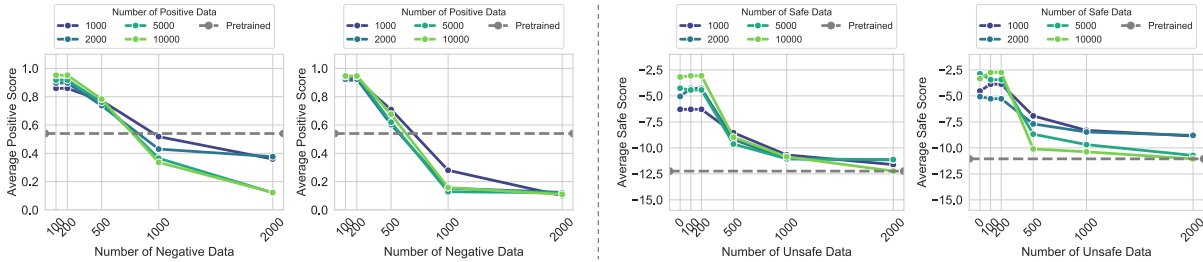


Figure 4: Experimental results for validating the existence of *rebound* (left: IMDb, right: Beavertails). The left part of each sub-figure is the performance of Gemma-2B while the right is Llama2-7B, respectively. Models trained with more *positive* data initially perform better but perform worse after fine-tuning with *negative* data.

the more *positive* the post-trained models’ performance, the more *negative* it becomes after *inverse alignment*. We validate tasks involving two opposing characteristics (e.g. safe and unsafe). We first train slices  $\{\theta_1, \theta_2, \dots, \theta_n\}$  based on a pre-trained model  $\theta_0$  using positive (e.g., safe) data of various volumes. Then we perform inverse finetuning on these slices using negative data (i.e., unsafe).

**Tasks and Datasets.** We select two tasks: positive generation and single-turn safe conversation. For the former, we use the data classified as positive or negative in the IMDb dataset (Maas et al., 2011). Referring to (Rafailov et al., 2024), we use the first 2-8 tokens of each complete text as a prompt for LLMs to generate the subsequent content. For the latter, we use the data classified as safe and unsafe in Beavertails (Ji et al., 2024c). We organized the positive sample sizes into  $\{1000, 2000, 5000, 10000\}$ , while the negative sample sizes, were divided into  $\{100, 200, 500, 1000, 2000\}$ .

**Evaluation and Metrics.** We collect the model’s responses on the reserved test prompts. Then we use score models provided by existing research to complete the performance evaluation. For positive style generation, we refer to (Rafailov et al., 2024)

and use the Sentiment Roberta model (Hartmann et al., 2023) to classify the responses, taking the proportion of all responses classified as positive as the model score. For single-turn safe dialogue, we use the cost model (Dai et al., 2024) to score the safety of each response, using the average score of all responses as the model score.

**Experiment Results.** We evaluate the *rebound* phenomenon on Llama2-7B (Touvron et al., 2023) and Gemma-2B (Team et al., 2024). The experimental results in Figure 4 show that, for models fine-tuned with a larger amount of positive sample data, their performance drops quicker under only a small amount of negative sample fine-tuning. Subsequently, the performance decline slows down and tends to stabilize. This result also confirms the previous conclusion: the initial rapid decline of model’s performance is due to *rebound*, as the model is far from the pre-trained distribution; while the later stabilization of the countermeasure is due to *resistance*, as the model is already close to the pre-trained distribution.

To assess the generalizability of the *rebound* phenomenon, we perform additional ablation studies focusing on alignment algorithms, evaluation

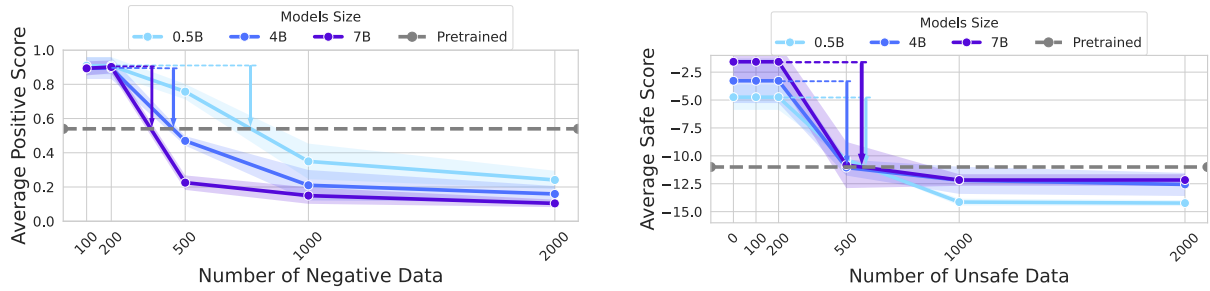


Figure 5: **Experimental results for validating *rebound* increases with model size (left: IMDb, right: Beavertails).** All single line covers positive data volume settings as Figure 4, with shadow denoting std. As the model size increases, the performance of the aligned model deteriorates more rapidly after fine-tuning with **negative** data.

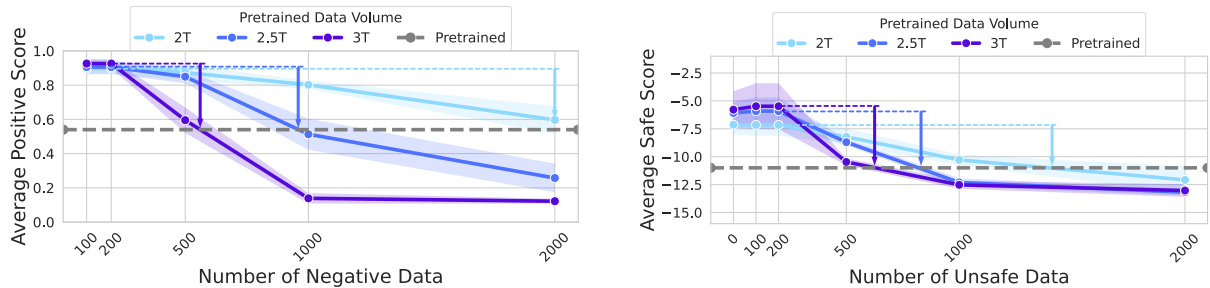


Figure 6: **Experimental results for validating *rebound* increases with pre-training data volume. (left: IMDb, right: Beavertails).** Each line covers positive data settings as Figure 4, with shadow denoting std. As pre-training data volume increases, aligned model performance deteriorates more rapidly after fine-tuning with **negative** data.

metrics, and fine-tuning directions. The results consistently validate the presence of the *rebound* phenomenon across language models. Further experimental details are available in Appendix B.2.

### 5.3 Internal Factor of *Rebound*

We conduct an in-depth analysis of the internal factors affecting the *rebound* phenomenon since it is crucial for the robustness of LLMs alignment. In particular, we investigate how model parameter scale and the amount of pre-training data influence *rebound*. Owing to space constraints, further experimental analyses are included in Appendix B.3.

***Rebound* Increases with Model Size.** To investigate how the *rebound* phenomenon varies with model size, we conducted experiments on Qwen models (Bai et al., 2023) with parameter scales of 0.5B, 4B, and 7B. The experimental results in Figure 5 show that as the model parameter size increases, the initial performance decline due to negative data fine-tuning is faster, while the subsequent decline is slower. This indicates that as the parameter size increases, there is an increase in *rebound* in response to both positive and negative data, further suggesting a positive correlation between model *elasticity* and parameter scale.

***Rebound* Increases with Pre-training Data Volume.** To verify that *rebound* increases with the growth of pre-training data, we vary pre-training slices (2.0T, 2.5T, and 3.0T) released by TinyLlama (Zhang et al., 2024a) and conduct the same experiment. As shown in Figure 6, when the pre-training data volume increases, the initial performance decline due to negative data fine-tuning is faster, while the subsequent decline is slower. It demonstrates that larger pre-training data volumes reinforce the *rebound* of LLMs, which is consistent with the inference proposed in Theorem 4.2.

## 6 Conclusion and Outlook

Our work uncovers a possible mechanism underlying the fragility of alignment: the *elasticity* of language models. We demonstrate that upon perturbation, the normalized compression rate of language models changes inversely with dataset size, making them more inclined to retain pre-training distributions while forgetting alignment distributions, thus resisting alignment. Through extensive experiments, we validate the universality of this *elasticity* effect, observing that the *elasticity* strengthens as model size and pre-training data volume scale up. Fundamentally, our findings support



a core assertion: language models exhibit *elasticity*, and thereby inherently resist alignment.

## 6.1 Limitations

Theory-wise, the primary limitation of our work is our specification of the *mass distribution* (Assumption A.7), and empirical studies on the exact form of this distribution shall be valuable. Experiment-wise, we have not systematically validated *elasticity* throughout the entire lifecycle of pre-training and alignment phases due to cost constraints. In future works, we plan to focus on whether this phenomenon is universally applicable, such as in multimodal models (Ji et al., 2024d; Huh et al., 2024). Additionally, we aim to theoretically uncover the relationship between model *elasticity* and *scaling laws* (Kaplan et al., 2020; Xiao et al., 2024), specifically determining the amount of training data required for *elasticity* to manifest and quantitatively analyzing whether *elasticity* intensifies as model parameters and pre-training data volume increase.

## 6.2 Broader Impacts

**Rethinking Fine-tuning and Model Weight from Resist Alignment.** Alignment fine-tuning methods aim to efficiently adjust the distribution of LLMs with minimal data to enhance their safety mechanisms. However, current alignment algorithms for language models may rely on optimization shortcuts, leading to local optima, while lacking involvement with the intrinsic mechanisms of the model (Qi et al., 2024a). From the perspective of language models’ *elasticity*, we require more robust alignment methods to ensure that modifications to model parameters extend beyond superficial changes (Qi et al., 2024a; Cohen et al., 2024), thereby preventing the emergence of more effective *inverse alignment* techniques that could give rise to extreme misalignment risks such as deceptive adversarial alignment (IDAIS-Beijing, 2024; Bengio et al., 2025; Marks et al., 2025; Baker et al., 2025) or alignment faking (Greenblatt et al., 2024). Although methods such as data cleansing during the training phase have proven to be an effective strategy for improving the malleability of a language model’s final distribution (He et al., 2024; Qi et al., 2024b), such approaches are not particularly cost-effective and feasible. In Appendix C.2, we provide a preliminary discussion of practical steps based on *elasticity* to mitigate *inverse alignment* risks, and we look forward to future research and practical developments grounded in *elasticity*, which

will enable the design of more robust alignment algorithms (Revel et al., 2024; Zhang et al., 2024b; Chen et al., 2024; Sheshadri et al., 2024; Liu et al., 2024; Li et al., 2025), ultimately achieving true and reliable alignment for LLMs.

**Rethinking Open-sourcing from Elasticity of LLMs.** Open-sourcing is a double-edged sword (Seger et al., 2023; Anwar et al., 2024; Kukreja et al., 2024). The public release of model weights enables the technical community to quickly identify and address potential vulnerabilities and conduct large-scale research on broad safety concerns (Touvron et al., 2023). These efforts collectively help sustain a balanced offence-defense dynamic (Jervis, 1978; Garfinkel and Dafoe, 2021) in the open-source environment and improve model security (Eiras et al., 2024). However, the risk of misuse associated with open-sourcing models is inevitable (Seger et al., 2023). Malicious fine-tuning of open-source models (Urbina et al., 2022; Sandbrink, 2023) and the use of the open-source models to facilitate system jailbreaks (Zou et al., 2023) can pose serious threats to public safety (Goldstein et al., 2023; Reuel et al., 2024). At present, the open-sourcing of powerful language models relies on safety alignment and rigorous security audits (Mökander et al., 2024) to ensure responsible use. However, if advanced *inverse alignment* techniques become feasible, even carefully aligned and audited open-source models could be reverted to their pre-alignment states at minimal cost, drastically lowering the barrier for jailbreaking aligned models and disrupting the offense-defense balance in the open-source ecosystem (Shapiro and Siegel, 2010). Therefore, it is crucial to develop alignment algorithms with fine-tuning robustness. We hope that deeper insights into the *elasticity* mechanisms of language models will drive the advancement of untunable alignment methods, enabling models to maintain reliable safety throughout their entire lifecycle (Madiaga, 2021).

## Acknowledgment

This work is sponsored by the National Natural Science Foundation of China (62376013, 623B2003, 624B100026). This work is supported by the Natural Science Foundation of Beijing (QY24041). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Usman Anwar, Abulhair Saparov, Javier Rando, Daniel Paleka, Miles Turpin, Peter Hase, Ekdeep Singh Lubana, Erik Jenner, Stephen Casper, Oliver Sourbut, et al. 2024. Foundational challenges in assuring alignment and safety of large language models. *arXiv preprint arXiv:2404.09932*.
- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022a. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022b. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. 2025. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. *arXiv preprint arXiv:2503.11926*.
- Yoshua Bengio, Sören Mindermann, Daniel Privitera, Tamay Besiroglu, Rishi Bommasani, Stephen Casper, Yejin Choi, Philip Fox, Ben Garfinkel, Danielle Goldfarb, et al. 2025. International ai safety report. *arXiv preprint arXiv:2501.17805*.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, Tony Tong Wang, Samuel Marks, Charbel-Raphael Segerie, Micah Carroll, Andi Peng, Phillip Christoffersen, Mehul Damani, Stewart Slocum, Usman Anwar, Anand Siththaranjan, Max Nadeau, Eric J Michaud, Jacob Pfau, Dmitrii Krasheninnikov, Xin Chen, Lauro Langosco, Peter Hase, Erdem Biyik, Anca Dragan, David Krueger, Dorsa Sadigh, and Dylan Hadfield-Menell. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *Transactions on Machine Learning Research*. Survey Certification.
- Yingfa Chen, Zhengyan Zhang, Xu Han, Chaojun Xiao, Zhiyuan Liu, Chen Chen, Kuai Li, Tao Yang, and Maosong Sun. 2024. Robust and scalable model editing for large language models. *arXiv preprint arXiv:2403.17431*.
- Michael K Cohen, Marcus Hutter, Yoshua Bengio, and Stuart Russell. 2024. RL, but don't do anything i wouldn't do. *arXiv preprint arXiv:2410.06213*.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, et al. 2024. Ultrafeedback: Boosting language models with scaled ai feedback. In *Forty-first International Conference on Machine Learning*.
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2024. Safe rlhf: Safe reinforcement learning from human feedback. In *The Twelfth International Conference on Learning Representations*.
- Grégoire Delétang, Anian Ruoss, Paul-Ambroise Duquenne, Elliot Catt, Tim Genewein, Christopher Mattern, Jordi Grau-Moya, Li Kevin Wenliang, Matthew Aitchison, Laurent Orseau, et al. 2023. Language modeling is compression. *arXiv preprint arXiv:2309.10668*.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, SHUM KaShun, and Tong Zhang. 2023. Raft: Reward ranked finetuning for generative foundation model alignment. *Transactions on Machine Learning Research*.
- Zhichen Dong, Zhanhui Zhou, Chao Yang, Jing Shao, and Yu Qiao. 2024. Attacks, defenses and evaluations for llm conversation safety: A survey. *arXiv preprint arXiv:2402.09283*.
- Francisco Eiras, Aleksandar Petrov, Bertie Vidgen, Christian Schroeder De Witt, Fabio Pizzati, Katherine Elkins, Supratik Mukhopadhyay, Adel Bibi, Botos Csaba, Fabro Steibel, et al. 2024. Near to mid-term risks and opportunities of open-source generative ai. *arXiv preprint arXiv:2404.17047*.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.
- Ben Garfinkel and Allan Dafoe. 2021. How does the offense-defense balance scale? In *Emerging Technologies and International Stability*, pages 247–274. Routledge.
- Josh A Goldstein, Girish Sastry, Micah Musser, Renee DiResta, Matthew Gentzel, and Katerina Sedova. 2023. Generative language models and automated influence operations: Emerging threats and potential mitigations. *arXiv preprint arXiv:2301.04246*.

- Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al. 2024. Alignment faking in large language models. *arXiv preprint arXiv:2412.14093*.
- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. 2023. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*.
- Jochen Hartmann, Mark Heitmann, Christian Siebert, and Christina Schamp. 2023. More than a feeling: Accuracy and application of sentiment analysis. *International Journal of Research in Marketing*, 40(1):75–87.
- Luxi He, Mengzhou Xia, and Peter Henderson. 2024. What’s in your "safe" data?: Identifying benign data that breaks safety. *arXiv preprint arXiv:2404.01099*.
- Robert Hooke. 2016. *Lectures de potentia restitutiva, or of spring explaining the power of springing bodies*. 6. John Martyn.
- Yuzhen Huang, Jinghan Zhang, Zifei Shan, and Junxian He. 2024. Compression represents intelligence linearly. *arXiv preprint arXiv:2404.09937*.
- Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Latham, Daniel M Ziegler, Tim Maxwell, Newton Cheng, et al. 2024. Sleeper agents: Training deceptive llms that persist through safety training. *arXiv preprint arXiv:2401.05566*.
- David A Huffman. 1952. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101.
- Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. 2024. Position: The platonic representation hypothesis. In *Forty-first International Conference on Machine Learning*.
- Marcus Hutter. 2005. *Universal Artificial Intelligence: Sequential Decisions Based On Algorithmic Probability*. Springer Science & Business Media.
- IDAIS-Beijing. 2024. [Consensus statement on red lines in artificial intelligence](#).
- Samyak Jain, Robert Kirk, Ekdeep Singh Lubana, Robert P Dick, Hidenori Tanaka, Edward Grefenstette, Tim Rocktäschel, and David Scott Krueger. 2023. Mechanistically analyzing the effects of fine-tuning on procedurally defined tasks. *arXiv preprint arXiv:2311.12786*.
- Robert Jervis. 1978. Cooperation under the security dilemma. *World politics*, 30(2):167–214.
- Jiaming Ji, Boyuan Chen, Hantao Lou, Donghai Hong, Borong Zhang, Xuehai Pan, Juntao Dai, and Yaodong Yang. 2024a. Aligner: Achieving efficient alignment through weak-to-strong correction. *arXiv preprint arXiv:2402.02416*.
- Jiaming Ji, Donghai Hong, Borong Zhang, Boyuan Chen, Josef Dai, Boren Zheng, Tianyi Qiu, Boxun Li, and Yaodong Yang. 2024b. Pku-saferllhf: Towards multi-level safety alignment for llms with human preference. *arXiv preprint arXiv:2406.15513*.
- Jiaming Ji, Mickel Liu, Josef Dai, Xuehai Pan, Chi Zhang, Ce Bian, Boyuan Chen, Ruiyang Sun, Yizhou Wang, and Yaodong Yang. 2024c. Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36.
- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, et al. 2023. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*.
- Jiaming Ji, Jiayi Zhou, Hantao Lou, Boyuan Chen, Donghai Hong, Xuyao Wang, Wenqi Chen, Kaile Wang, Rui Pan, Jiahao Li, et al. 2024d. Align anything: Training all-modality models to follow instructions with language feedback. *arXiv preprint arXiv:2412.15838*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Sanjay Kukreja, Tarun Kumar, Amit Purohit, Abhijit Dasgupta, and Debashis Guha. 2024. A literature survey on open source large language models. In *Proceedings of the 2024 7th International Conference on Computers in Management and Business*, pages 133–143.
- Leon Lang, Davis Foote, Stuart J Russell, Anca Dragan, Erik Jenner, and Scott Emmons. 2024. When your ais deceive you: Challenges of partial observability in reinforcement learning from human feedback. *Advances in Neural Information Processing Systems*, 37:93240–93299.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbone, and Abhinav Rastogi. 2023. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*.
- Bohan Li, Longxu Dou, Yutai Hou, Yunlong Feng, Honglin Mu, Enbo Wang, Qingfu Zhu, Qinghua Sun, and Wanxiang Che. 2025. Mixpro: Simple yet effective data augmentation for prompt-based learning. *International Journal of Machine Learning and Cybernetics*, pages 1–20.

- Ziniu Li, Tian Xu, Yushun Zhang, Yang Yu, Ruoyu Sun, and Zhi-Quan Luo. 2023. Remax: A simple, effective, and efficient method for aligning large language models. *arXiv preprint arXiv:2310.10505*.
- Lizhi Lin, Honglin Mu, Zenan Zhai, Minghan Wang, Yuxia Wang, Renxi Wang, Junjie Gao, Yixuan Zhang, Wanxiang Che, Timothy Baldwin, et al. 2025. Against the achilles' heel: A survey on red teaming for generative models. *Journal of Artificial Intelligence Research*, 82:687–775.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252.
- Tianqi Liu, Wei Xiong, Jie Ren, Lichang Chen, Junru Wu, Rishabh Joshi, Yang Gao, Jiaming Shen, Zhen Qin, Tianhe Yu, et al. 2024. Rrm: Robust reward model training mitigates reward hacking. *arXiv preprint arXiv:2409.13156*.
- Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150.
- Tambiama Madiaga. 2021. Artificial intelligence act. *European Parliament: European Parliamentary Research Service*.
- Samuel Marks, Johannes Treutlein, Trenton Bricken, Jack Lindsey, Jonathan Marcus, Siddharth Mishra-Sharma, Daniel Ziegler, Emmanuel Ameisen, Joshua Batson, Tim Belonax, et al. 2025. Auditing language models for hidden objectives. *arXiv preprint arXiv:2503.10965*.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235.
- Chris Mingard, Guillermo Valle-Pérez, Joar Skalse, and Ard A Louis. 2021. Is sgd a bayesian sampler? well, almost. *Journal of Machine Learning Research*, 22(79):1–64.
- Jakob Mökander, Jonas Schuett, Hannah Rose Kirk, and Luciano Floridi. 2024. Auditing large language models: a three-layered approach. *AI and Ethics*, 4(4):1085–1115.
- Kevin P Murphy. 2023. *Probabilistic machine learning: Advanced topics*. MIT press.
- Mark EJ Newman. 2005. Power laws, pareto distributions and zipf's law. *Contemporary physics*, 46(5):323–351.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Peter S Park, Simon Goldstein, Aidan O'Gara, Michael Chen, and Dan Hendrycks. 2024. Ai deception: A survey of examples, risks, and potential solutions. *Patterns*, 5(5).
- Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek Mittal, and Peter Henderson. 2024a. Safety alignment should be made more than just a few tokens deep. *arXiv preprint arXiv:2406.05946*.
- Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2024b. Fine-tuning aligned language models compromises safety, even when users do not intend to! In *The Twelfth International Conference on Learning Representations*.
- Chen Qian, Jie Zhang, Wei Yao, Dongrui Liu, Zhenfei Yin, Yu Qiao, Yong Liu, and Jing Shao. 2024. Towards tracing trustworthiness dynamics: Revisiting pre-training period of large language models. *arXiv preprint arXiv:2402.19465*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Anka Reuel, Ben Bucknall, Stephen Casper, Tim Fist, Lisa Soder, Onni Aarne, Lewis Hammond, Lujain Ibrahim, Alan Chan, Peter Wills, et al. 2024. Open problems in technical ai governance. *arXiv preprint arXiv:2407.14981*.
- Manon Revel, Matteo Cargnelutti, Tyna Eloundou, and Greg Leppert. 2024. Seal: Systematic error analysis for value alignment. *arXiv preprint arXiv:2408.10270*.
- Jonas B Sandbrink. 2023. Artificial intelligence and biological misuse: Differentiating risks of language models and biological design tools. *arXiv preprint arXiv:2306.13952*.
- Elizabeth Seger, Noemi Dreksler, Richard Moulange, Emily Dardaman, Jonas Schuett, K Wei, Christoph Winter, Mackenzie Arnold, Seán Ó hÉigearthaigh, Anton Korinek, et al. 2023. Open-sourcing highly capable foundation models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives. *arXiv preprint arXiv:2311.09227*.
- Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423.

- Jacob N Shapiro and David A Siegel. 2010. Is this paper dangerous? balancing secrecy and openness in counterterrorism. *Security Studies*, 19(1):66–98.
- Abhay Sheshadri, Aidan Ewart, Phillip Guo, Aengus Lynch, Cindy Wu, Vivek Hebbar, Henry Sleight, Asa Cooper Stickland, Ethan Perez, Dylan Hadfield-Menell, et al. 2024. Targeted latent adversarial training improves robustness to persistent harmful behaviors in llms. *arXiv e-prints*, pages arXiv–2407.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Fabio Urbina, Filippa Lentzos, Cédric Invernizzi, and Sean Ekins. 2022. Dual use of artificial-intelligence-powered drug discovery. *Nature machine intelligence*, 4(3):189–191.
- Boyi Wei, Kaixuan Huang, Yangsibo Huang, Tinghao Xie, Xiangyu Qi, Mengzhou Xia, Prateek Mittal, Mengdi Wang, and Peter Henderson. 2024. Assessing the brittleness of safety alignment via pruning and low-rank modifications. *arXiv preprint arXiv:2402.05162*.
- Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. 2016. A survey of transfer learning. *Journal of Big data*, 3:1–40.
- Jiaxin Wen, Ruiqi Zhong, Akbir Khan, Ethan Perez, Jacob Steinhardt, Minlie Huang, Samuel R Bowman, He He, and Shi Feng. 2024. Language models learn to mislead humans via rlhf. *arXiv preprint arXiv:2409.12822*.
- Ian H Witten, Radford M Neal, and John G Cleary. 1987. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540.
- Chaojun Xiao, Jie Cai, Weilin Zhao, Guoyang Zeng, Biyuan Lin, Jie Zhou, Zhi Zheng, Xu Han, Zhiyuan Liu, and Maosong Sun. 2024. Densing law of llms. *arXiv preprint arXiv:2412.04315*.
- Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang. 2024. Iterative preference learning from human feedback: Bridging theory and practice for rlhf under kl-constraint. In *Forty-first International Conference on Machine Learning*.
- Kevin Yang, Dan Klein, Asli Celikyilmaz, Nanyun Peng, and Yuandong Tian. 2023a. Rlcd: Reinforcement learning from contrastive distillation for language model alignment. *arXiv preprint arXiv:2307.12950*.
- Xianjun Yang, Xiao Wang, Qi Zhang, Linda Petzold, William Yang Wang, Xun Zhao, and Dahua Lin. 2023b. Shadow alignment: The ease of subverting safely-aligned language models. *arXiv preprint arXiv:2310.02949*.
- Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. 2024. Alignment for honesty. *Advances in Neural Information Processing Systems*, 37:63565–63598.
- Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024a. Tinyllama: An open-source small language model. *arXiv preprint arXiv:2401.02385*.
- Xuanliang Zhang, Dingzirui Wang, Longxu Dou, Qingfu Zhu, and Wanxiang Che. 2025. A survey of table reasoning with large language models. *Frontiers of Computer Science*, 19(9):199348.
- Zhexin Zhang, Junxiao Yang, Pei Ke, Shiyao Cui, Chujie Zheng, Hongning Wang, and Minlie Huang. 2024b. Safe unlearning: A surprisingly effective and generalizable solution to defend against jailbreak attacks. *arXiv preprint arXiv:2407.02855*.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.
- Jiayi Zhou, Jiaming Ji, Josef Dai, and Yaodong Yang. 2025. Sequence to sequence reward modeling: Improving rlhf by language feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 27765–27773.
- Zhanhui Zhou, Jie Liu, Jing Shao, Xiangyu Yue, Chao Yang, Wanli Ouyang, and Yu Qiao. 2023. Beyond one-preference-fits-all alignment: Multi-objective direct preference optimization. *arXiv preprint arXiv:2310.03708*.
- Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. 2020. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76.
- George Kingsley Zipf. 1946. The psychology of language. In *Encyclopedia of psychology*, pages 332–341. Philosophical Library.
- Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.

## A Assumptions and Proofs

**Assumption A.1** (Binary Tokens). For simplicity and without loss of generality, we assume that all datasets share a uniform token table containing only binary tokens (specifically  $\emptyset/1$ ).

**Theorem A.2** (Ideal Code Length). Consider a finite parameter model  $p_\theta(\cdot)$  training on dataset  $\mathcal{D}$ , the ideal code length  $\mathcal{L}_{p_\theta}(\mathbf{x})$  of a random response  $\mathbf{x}$  compressed by  $p_\theta$  can be expressed as follows:

$$\mathbb{E}[\mathcal{L}_{p_\theta}(\mathbf{x})] = \left\lceil \frac{|\mathbf{x}|}{d} \right\rceil \left[ - \sum_{l=1}^d \sum_{j=1}^{2^{l-1}} p_{lj} \log p_{lj} \right],$$

where  $d$  represents the depth of the  $\mathcal{T}_\mathcal{D}$  after pruning under Definition 3.3 protocol,  $p_{lj}$  represents the probability values of the EOS nodes for the  $j$ -th node at the  $l$ -th layer.

*Proof.* When  $|\mathbf{x}| \leq d$ , the compression protocol defined in Definition 3.3 can perfectly compress  $\mathbf{x}$ . Hence, the expectation of the ideal code length  $\mathcal{L}_{p_\theta}(\mathbf{x})$  satisfies:

$$\mathbb{E}[\mathcal{L}_{p_\theta}(\mathbf{x})] = \left[ - \sum_{l=1}^d \sum_{j=1}^{2^{l-1}} p_{lj} \log p_{lj} \right],$$

where  $d$  represents the depth of the pruned tree  $\mathcal{T}'_\mathcal{D}$  and  $p_{lj}$  represents the probability values of the EOS nodes for the  $j$ -th node at the  $l$ -th layer.

Now consider  $sd \leq |\mathbf{x}| \leq (s+1)d$ . Let us suppose that  $\mathbf{x} = (\mathbf{x}_1 \cdots \mathbf{x}_s \mathbf{x}_{s+1})$ , where  $|\mathbf{x}_k| = d$ , for  $k \in \{1, \dots, s\}$  and  $|\mathbf{x}_{s+1}| \leq d$ . In this case,  $\mathbf{x}$  cannot be perfectly compressed by the model. Hence, the compression of  $\mathbf{x}$  needs to be performed in segments, and the length of each segment is not greater than  $d$ .

$$\begin{aligned} \mathbb{E}_\mathbf{x}[\mathcal{L}_{p_\theta}(\mathbf{x})] &= \mathbb{E}_{\mathbf{x}_1}[\mathcal{L}_{p_\theta}(\mathbf{x}_1)] + \mathbb{E}_{\mathbf{x}_1} \mathbb{E}_{(\mathbf{x}_2 \dots \mathbf{x}_{s+1})}[\mathcal{L}_{p_\theta}((\mathbf{x}_2 \dots \mathbf{x}_{s+1}) | \mathbf{x}_1)] \\ &= \mathbb{E}_{\mathbf{x}_1}[\mathcal{L}_{p_\theta}(\mathbf{x}_1)] + \sum_{\mathbf{x}_1} p(\mathbf{x}_1) \sum_{(\mathbf{x}_2 \dots \mathbf{x}_{s+1})} p(\mathbf{x}_2 \dots \mathbf{x}_{s+1} | \mathbf{x}_1) \cdot \mathcal{L}_{p_\theta}((\mathbf{x}_2 \dots \mathbf{x}_{s+1})) \\ &= \mathbb{E}_{\mathbf{x}_1}[\mathcal{L}_{p_\theta}(\mathbf{x}_1)] + \mathbb{E}_{(\mathbf{x}_2 \dots \mathbf{x}_{s+1})}[\mathcal{L}_{p_\theta}((\mathbf{x}_2 \dots \mathbf{x}_{s+1}))] \\ &= \sum_{k=1}^{s+1} \mathbb{E}_{\mathbf{x}_k}[\mathcal{L}_{p_\theta}(\mathbf{x}_k)] \\ &= \left\lceil \frac{|\mathbf{x}|}{d} \right\rceil \left[ - \sum_{l=1}^d \sum_{j=1}^{2^{l-1}} p_{lj} \log p_{lj} \right], \end{aligned}$$

thus the proof is completed.  $\square$

**Definition A.3** (Mass Distribution in Token Tree). Consider the sample space  $\Omega$  consisting of all responses in dataset  $\mathcal{D}$ . The probability distribution  $\mathcal{P}_\mathcal{D}$  of all subtrees at the  $d$ -th level nodes of  $\mathcal{T}_\mathcal{D}$  is a mapping from  $\Omega$  to  $[0, 1]$ . Let  $X_\mathcal{D}$  be the random variable representing the probability value taken at each leaf multiplied by the number of leaves. The mass distribution  $\mathcal{P}_{mass}$  represents the probability that  $X_\mathcal{D}$  takes the corresponding probability value. According to the definition of  $\mathcal{P}_{mass}$ ,  $\mathbb{E}[X_\mathcal{D}] = 1$ .

*Remark A.4* (Mixture of Mass Distribution). For independently and differently distributed datasets  $\mathcal{D}_1, \dots, \mathcal{D}_N$ ,  $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$  is a mixture of these datasets. For the pruned trees  $\mathcal{T}_1, \dots, \mathcal{T}_N$  of these datasets with depth  $d$ , the random variables of their leaf nodes satisfy the following relationship:

$$X_\mathcal{D} = \frac{\sum_{k=1}^N |\mathcal{D}_k| X_{\mathcal{D}_k}}{\sum_{k=1}^N |\mathcal{D}_k|},$$

where  $X_{\mathcal{D}_k}$  follows the mass distribution  $\mathcal{P}_{mass}^k$ . For  $X_{\mathcal{D}_{k_i}}$  and  $X_{\mathcal{D}_{k_j}}$  from different datasets,  $X_{\mathcal{D}_{k_i}}$  and  $X_{\mathcal{D}_{k_j}}$  are independent of each other.

**Lemma A.5** (Entropy of Mass Distribution). Consider the pruned trees  $\mathcal{T}'_k$  and  $\mathcal{T}'$  of dataset  $\mathcal{D}_k$  and  $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$ . Denote that the response distribution and the mass distribution of  $\mathcal{T}'$ 's tree nodes are  $\mathcal{P}^{\mathcal{D}}$  and  $\mathcal{P}_{mass}$ . Similarly, we define the response distribution  $\mathcal{P}_k^{\mathcal{D}}$  and the mass distribution  $\mathcal{P}_{mass}^k$ . When the number of the leaf nodes  $M$  is sufficiently large, the cross-entropy of the response distribution can be rewritten as follows.

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{P}_k} [-p^{\mathcal{D}_k} \log p^{\mathcal{D}}] = \mathbb{E}_{X_{\mathcal{D}_k} \sim \mathcal{P}_{mass}^k, X_{\mathcal{D}} \sim \mathcal{P}_{mass}} [-X_{\mathcal{D}_k} \log X_{\mathcal{D}}] + \log M,$$

where  $p^{\mathcal{D}}$ ,  $p^{\mathcal{D}_k}$  stand for the probability of the leaf nodes of  $\mathcal{T}'$ ,  $\mathcal{T}'_k$  while  $X_{\mathcal{D}_k}$ ,  $X_{\mathcal{D}}$  stand for the random variables of the probability of the leaf nodes in  $\mathcal{T}'$ ,  $\mathcal{T}'_k$ , respectively.

*Proof.* Let  $M$  be the number of leaf nodes of  $\mathcal{T}'$ . According to the definitions of the response distribution  $\mathcal{P}$  and mass distribution  $\mathcal{P}_{mass}$ , we have  $Mp^{\mathcal{D}_j} = X_{\mathcal{D}_j}$ ,  $\forall j \in \{1, \dots, N\}$ . Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_k} [-p^{\mathcal{D}_k} \log p^{\mathcal{D}}] &= \sum_{i=1}^M -p_i^{\mathcal{D}_k} \log p_i^{\mathcal{D}} \\ &= \sum_{i=1}^M -\frac{1}{M} X_{i, \mathcal{D}_k} \log X_{i, \mathcal{D}} + \log M \\ &= \mathbb{E}_{X_{\mathcal{D}_k} \sim \mathcal{P}_{mass}^k, X_{\mathcal{D}} \sim \mathcal{P}_{mass}} [-X_{\mathcal{D}_k} \log X_{\mathcal{D}}] + \log M. \end{aligned}$$

□

*Remark A.6.* In Lemma A.5,  $X_{\mathcal{D}_k}$  are assumed to be independent. However due to  $\sum_{i=1}^M p_i^{\mathcal{D}_k} = 1$ , the  $X_{\mathcal{D}_k}$  are not actually independent. Considering that  $M$  is sufficiently large in our subsequent analysis, we can regard the independence of  $X_{\mathcal{D}_k}$  as a good approximation.

We assume that the mass distribution of the segment follows a heavy-tailed Pareto distribution, with supporting evidence from [Mingard et al. \(2021\)](#); [Zipf \(1946\)](#).

**Assumption A.7** (The Pareto Distribution). We assume that the mass distribution of pruned token trees  $\mathcal{T}$  of depth  $d$ , across different datasets, follows a Pareto distribution with identical parameters:

$$p_X(x) = \begin{cases} \frac{\alpha c^\alpha}{x^{\alpha+1}} & x \geq c, \\ 0 & x < c, \end{cases}$$

where  $\alpha, c$  are parameters of the Pareto distribution.

**Theorem A.8** (Elasticity of Language Models). Consider the pre-training dataset  $\mathcal{D}_p = \bigcup_{i=1}^3 \mathcal{D}_i$ , the alignment dataset  $\mathcal{D}_a$ , and the perturbation dataset  $\mathcal{D}_t$ , with the model  $p_\theta(\cdot)$  trained on  $\mathcal{D} = \mathcal{D}_p \cup \mathcal{D}_a \cup \mathcal{D}_t$ . Assume that  $\mathcal{D}_a \stackrel{d}{\sim} \mathcal{D}_2$ ,  $\mathcal{D}_t \stackrel{d}{\sim} \mathcal{D}_3$ , and  $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3$  are each distributed according to a Pareto mass distribution ([Newman, 2005](#)). As the perturbation data volume  $|\mathcal{D}_t|$  varies,  $\gamma_{p_\theta}^{\mathcal{D}_p/\mathcal{D}}$  and  $\gamma_{p_\theta}^{\mathcal{D}_a/\mathcal{D}}$  satisfy:

$$\begin{aligned} \frac{d\gamma_{p_\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} &= \Theta \left( -k \frac{d\gamma_{p_\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} \right), \\ \frac{d\gamma_{p_\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} &< 0, \quad \frac{d\gamma_{p_\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} > 0, \end{aligned} \tag{4}$$

where  $l = \frac{|\mathcal{D}_t|}{|\mathcal{D}_a|} \ll 1$ ,  $k = \frac{|\mathcal{D}_p|}{|\mathcal{D}_a|} \gg 1$ , and  $\{\mathcal{D}_i\}_{i=1}^3$  are datasets of equal cardinality.

*Proof.* For the sake of convenience in calculations, we first use Lemma A.5 to replace the Shannon entropy of response distribution.

$$\begin{aligned} \frac{d\gamma_{p\theta}^{\mathcal{D}_j/\mathcal{D}}}{dl} &= \frac{d(\mathbb{E}_{\mathbf{x}\sim\mathcal{P}_j}[-p^{\mathcal{D}_j} \log p^{\mathcal{D}}] - \log M)}{dl} \\ &= \frac{d\left(\mathbb{E}_{X_{\mathcal{D}_j}\sim\mathcal{P}_{mass}^j, X_{\mathcal{D}}\sim\mathcal{P}_{mass}}[-X_{\mathcal{D}_j} \log X_{\mathcal{D}}]\right)}{dl}. \end{aligned}$$

According to Assumption A.7,  $X_{\mathcal{D}_j}$  follows a Pareto distribution with the parameters  $\alpha$  and  $c$ . Hence,

$$\begin{aligned} &\mathbb{E}_{X_{\mathcal{D}_p}\sim\mathcal{P}_{mass}^p, X_{\mathcal{D}}\sim\mathcal{P}_{mass}}[-X_{\mathcal{D}_p} \log X_{\mathcal{D}}] \\ &= -\int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha}}{\prod_{i=1}^3 x_i^{\alpha+1}} \cdot \frac{x_1 + x_2 + x_3}{3} \cdot \log \frac{\sum_{i\in\{p,a,t\}} |\mathcal{D}_i| x_i}{\sum_{i\in\{p,a,t\}} |\mathcal{D}_i|} dx_1 dx_2 dx_3 \\ &= -\int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha} (x_1 + x_2 + x_3)}{3 \prod_{i=1}^3 x_i^{\alpha+1}} \log \frac{\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3}{k+l+1} dx_1 dx_2 dx_3 \\ &\mathbb{E}_{X_{\mathcal{D}_a}\sim\mathcal{P}_{mass}^a, X_{\mathcal{D}}\sim\mathcal{P}_{mass}}[-X_{\mathcal{D}_a} \log X_{\mathcal{D}}] \\ &= -\int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha} x_2}{\prod_{i=1}^3 x_i^{\alpha+1}} \log \frac{\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3}{k+l+1} dx_1 dx_2 dx_3 \end{aligned}$$

where  $j = 1, 2, 3$ . Therefore,  $\frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl}$  and  $\frac{d\gamma_{p\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl}$  can be written as:

$$\frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} = -\frac{dS_1}{dl}, \quad \frac{d\gamma_{p\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} = -\frac{dS_2}{dl} \quad (5)$$

where

$$\begin{aligned} S_1 &= \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha} (x_1 + x_2 + x_3)}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \log \frac{\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3}{k+l+1} dx_1 dx_2 dx_3 \\ S_2 &= \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha}}{x_1^{\alpha+1} x_2^{\alpha} x_3^{\alpha+1}} \log \frac{\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3}{k+l+1} dx_1 dx_2 dx_3. \end{aligned}$$

Proving that  $\frac{d\gamma_{p\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} = \Theta\left(-k \frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl}\right)$  is equivalent to proving:

$$\lim_{k\rightarrow+\infty, l\rightarrow 0} \frac{k \cdot \frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} + \frac{d\gamma_{p\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl}}{k \cdot \frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl}} = 0, \quad (6)$$

By substituting (5) into (6), we have

$$\lim_{k\rightarrow+\infty, l\rightarrow 0} \frac{k \cdot \frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} + \frac{d\gamma_{p\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl}}{k \cdot \frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl}} = \frac{\lim_{k\rightarrow+\infty, l\rightarrow 0} \left(\frac{k \cdot dS_1 + dS_2}{dl}\right)}{\lim_{k\rightarrow+\infty, l\rightarrow 0} \frac{k \cdot dS_1}{dl}}. \quad (7)$$



Now calculate the values of  $\frac{dS_1}{dl}$  for the case when  $l \rightarrow 0$  and  $k$  is sufficiently large.

$$\begin{aligned}
& \lim_{l \rightarrow 0} \frac{dS_1}{dl} \\
&= \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{(x_1 + x_2 + x_3)}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \\
& \quad \cdot \frac{(k+1+l)x_3 - (\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3)}{(k+1+l)(\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3)} dx_1 dx_2 dx_3 \\
&= \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{(x_1 + x_2 + x_3)x_3}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3)} dx_1 dx_2 dx_3 \\
& \quad - \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{(x_1 + x_2 + x_3)}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (k+1+l)} dx_1 dx_2 dx_3 \\
&= \alpha^3 c^{3\alpha} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \lim_{l \rightarrow 0} \frac{(x_1 + x_2 + x_3)x_3}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3)} dx_1 dx_2 dx_3 \\
& \quad - \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (k+1+l)} dx_1 dx_2 dx_3 \\
&= \alpha^3 c^{3\alpha} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1 + x_2 + x_3}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \\
& \quad \cdot \lim_{l \rightarrow 0} \left[ \frac{x_3}{\frac{k}{3}(x_1 + x_2 + x_3) + x_2} - \frac{x_3^2}{[\frac{k}{3}(x_1 + x_2 + x_3) + x_2]^2} l \right] dx_1 dx_2 dx_3 \\
& \quad - \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (k+1+l)} dx_1 dx_2 dx_3 \tag{8}
\end{aligned}$$

$$\begin{aligned}
&= \alpha^3 c^{3\alpha} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1 + x_2 + x_3}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \frac{x_3}{\frac{k}{3}(x_1 + x_2 + x_3) + x_2} dx_1 dx_2 dx_3 \\
& \quad - \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (k+1+l)} dx_1 dx_2 dx_3 \\
&= \alpha^3 c^{3\alpha} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1 + x_2 + x_3}{3x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \left[ \frac{x_3}{(x_1 + x_2 + x_3)k} - \frac{x_2 x_3}{(x_1 + x_2 + x_3)^2 k^2} \frac{9}{k} \right] dx_1 dx_2 dx_3 \\
& \quad - \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1} (k+1+l)} dx_1 dx_2 dx_3 \tag{9}
\end{aligned}$$

$$\begin{aligned}
&= \alpha^3 c^{3\alpha} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \frac{1}{k(k+1)} dx_1 dx_2 dx_3 \\
& \quad - \alpha^3 c^{3\alpha} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \frac{3x_2 x_3}{k^2(x_1 + x_2 + x_3)} dx_1 dx_2 dx_3 \tag{10}
\end{aligned}$$

Here, we apply a Taylor expansion to Equation (8) and Equation (9) with respect to  $l$  and  $\frac{1}{k}$ , respectively.

Similarly, we can calculate the value of  $\frac{dS_2}{dl}$ .

$$\begin{aligned}
& \lim_{l \rightarrow 0} \frac{dS_2}{dl} \\
&= \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_2}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \\
& \quad \cdot \frac{(k+1+l)x_3 - (\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3)}{(k+1+l)(\frac{k}{3}(x_1 + x_2 + x_3) + x_2 + lx_3)} dx_1 dx_2 dx_3 \\
&= \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha}}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \cdot \left( \frac{3x_2 x_3}{k(x_1 + x_2 + x_3)} - \frac{1}{k+1+l} \right) dx_1 dx_2 dx_3 \tag{11}
\end{aligned}$$

As a result, Equation (7) can be written as:

$$\begin{aligned}
& \lim_{k \rightarrow +\infty, l \rightarrow 0} k \cdot \frac{d\gamma_{p\theta}^{\mathcal{D}_p/\mathcal{D}}}{dl} + \frac{d\gamma_{p\theta}^{\mathcal{D}_a/\mathcal{D}}}{dl} \\
&= \lim_{k \rightarrow +\infty, l \rightarrow 0} \left( \frac{k \cdot dS_1 + dS_2}{dl} \right) \\
&= \frac{\lim_{k \rightarrow +\infty, l \rightarrow 0} \frac{k \cdot dS_1}{dl}}{\lim_{k \rightarrow +\infty, l \rightarrow 0} \left( k \cdot \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha}}{\prod_{i=1}^3 x_i^{\alpha+1}} \left[ \frac{x_1}{k(k+1)} - \frac{3x_2x_3}{k^2(\sum_{i=1}^3 x_i)} \right] dx_1 dx_2 dx_3 \right. \\
&\quad \left. + \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha}}{\prod_{i=1}^3 x_i^{\alpha+1}} \cdot \left( \frac{3x_2x_3}{k(\sum_{i=1}^3 x_i)} - \frac{1}{k+1+l} \right) dx_1 dx_2 dx_3 \right)} \\
&= \frac{\lim_{k \rightarrow +\infty, l \rightarrow 0} k \cdot \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{1}{\prod_{i=1}^3 x_i^{\alpha+1}} \left[ \frac{x_1}{k(k+1)} - \frac{3x_2x_3}{k^2(\sum_{i=1}^3 x_i)} \right] dx_1 dx_2 dx_3}{\lim_{k \rightarrow +\infty, l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\alpha^3 c^{3\alpha}}{\prod_{i=1}^3 x_i^{\alpha+1}} \cdot \frac{l}{(k+1)(k+1+l)}} \\
&= \frac{\lim_{k \rightarrow +\infty, l \rightarrow 0} k \cdot \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{1}{\prod_{i=1}^3 x_i^{\alpha+1}} \left[ \frac{x_1}{k(k+1)} - \frac{3x_2x_3}{k^2(\sum_{i=1}^3 x_i)} \right] dx_1 dx_2 dx_3}{=0}.
\end{aligned} \tag{12}$$

Next, we prove that  $\frac{d\gamma_{p\theta}^{\mathcal{D}_1/\mathcal{D}}}{dl} < 0$ . For sufficiently large  $k$ , since

$$\begin{aligned}
& \lim_{l \rightarrow 0} \frac{dS_1}{dl} \\
&= \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{x_1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \frac{1}{k(k+1)} dx_1 dx_2 dx_3 \\
&\quad - \alpha^3 c^{3\alpha} \lim_{l \rightarrow 0} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{1}{x_1^{\alpha+1} x_2^{\alpha+1} x_3^{\alpha+1}} \frac{3x_2x_3}{k^2(x_1+x_2+x_3)} dx_1 dx_2 dx_3 \\
&= \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{\frac{k}{k+1} x_1(x_1+x_2+x_3) - 3x_1x_3}{k(x_1+x_2+x_3)x_1^{\alpha+1}x_2^{\alpha+1}x_3^{\alpha+1}} dx_1 dx_2 dx_3 \\
&:= \frac{1}{k} \int_c^{+\infty} \int_c^{+\infty} \int_c^{+\infty} \frac{f(x_1, x_2, x_3)}{g(x_1, x_2, x_3)} dx_1 dx_2 dx_3.
\end{aligned} \tag{13}$$

For  $c < u < v < w$ , let

$$\begin{aligned}
\hat{f} &= x_1(x_1+x_2+x_3) - 3x_1x_3 = x_1(x_1+x_2-2x_3), \\
h(u, v, w) &= \sum_{\pi \in S_3} f(\pi(u, v, w)), \\
\hat{h}(u, v, w) &= \sum_{\pi \in S_3} \hat{f}(\pi(u, v, w)).
\end{aligned} \tag{14}$$

Here,  $S_3$  denotes the symmetric group of degree 3 and  $\pi \in S_3$  denotes a permutation in  $S_3$ . Observe that

$$\begin{aligned}
& \lim_{k \rightarrow +\infty} h(u, v, w) = \hat{h}(u, v, w)^-, \\
& \hat{h}(u, v, w) = \sum_{\pi \in S_3} \hat{f}(\pi(u, v, w)) \\
&= \sum_{\substack{\pi \in S_3 \\ (u_0, v_0, w_0) = \pi(u, v, w)}} u_0(u_0 + v_0 - 2w_0) \\
&= \sum_{\substack{\pi \in S_3 \\ (u_0, v_0, w_0) = \pi(u, v, w)}} \frac{1}{2}(u_0 - v_0)^2 > 0.
\end{aligned} \tag{15}$$

Hence, we have

$$\lim_{k \rightarrow +\infty} h(u, v, w) > 0. \quad (16)$$

Therefore,

$$\lim_{k \rightarrow +\infty} \frac{h(u, v, w)}{g(u, v, w)} = \lim_{k \rightarrow +\infty} \frac{\Theta(1)}{\Theta(k)} = 0^+. \quad (17)$$

Therefore,

$$\begin{aligned} \lim_{k \rightarrow +\infty, l \rightarrow 0} \frac{dS_1}{dl} &= \lim_{k \rightarrow +\infty} \iiint_{(u,v,w) \in (c, +\infty)^3: u < v < w} \frac{h(u, v, w)}{g(u, v, w)} dudvdw \\ &= \iiint_{(u,v,w) \in (c, +\infty)^3: u < v < w} \lim_{k \rightarrow +\infty, l \rightarrow 0} \frac{h(u, v, w)}{g(u, v, w)} dudvdw \\ &= \iiint_{(u,v,w) \in (c, +\infty)^3: u < v < w} \lim_{k \rightarrow +\infty, l \rightarrow 0} 0^+ dudvdw \\ &= 0^+. \end{aligned} \quad (18)$$

Hence, for sufficiently large  $k$ , we have

$$\lim_{l \rightarrow 0} \frac{d\gamma_{p\theta}^{\mathcal{D}_1/\mathcal{D}}}{dl} = \lim_{l \rightarrow 0} -\frac{dS_1}{dl} < 0. \quad (19)$$

Thus, the proof is completed.  $\square$

## B Additional Experiment Results

### B.1 Additional Experiment of Language Models' Resistance

To broadly validate the phenomenon of *Resistance* in language models, we extend the experimental setup described in Section 5.1. Specifically, we use Llama2-7B (Touvron et al., 2023) as the base model and perform a finer-grained snapshot division within *forward alignment*, covering a wider range of  $k$  and  $l$ . As the evaluation metric, we measure the change in relative KL divergence between the distributions of model  $\theta_l$  and  $\theta_k$ , obtained by applying *forward alignment* and *inverse alignment* to slices  $\theta_k$  and  $\theta_l$ , respectively, relative to the original distributions of  $\theta_l$  and  $\theta_k$ . The experimental results, which is shown in the Table 2, demonstrate that the KL divergence under *inverse alignment* is substantially smaller than that under *forward alignment*, consistent with the conclusions presented in the main text.

Training Steps	$\theta_1$ vs. $\theta_2$	$\theta_1$ vs. $\theta_3$	$\theta_1$ vs. $\theta_4$	$\theta_1$ vs. $\theta_5$	$\theta_1$ vs. $\theta_6$	$\theta_2$ vs. $\theta_3$	$\theta_2$ vs. $\theta_4$
<i>Forward Alignment</i>	0.4568 $\uparrow$	0.4932 $\uparrow$	0.5929 $\uparrow$	0.4439 $\uparrow$	0.4109 $\uparrow$	0.3778 $\uparrow$	0.5580 $\uparrow$
<i>Inverse Alignment</i>	0.2796 $\downarrow$	0.2863 $\downarrow$	0.3960 $\downarrow$	0.3995 $\downarrow$	0.3412 $\downarrow$	0.2752 $\downarrow$	0.3742 $\downarrow$

Table 2: **KL Divergence between different snapshots.** The results show that *forward alignment* exhibits higher values compared to *inverse alignment*, providing evidence for the existence of *Resistance*.

### B.2 Additional Experiment of Language Models' Rebound

**Ablations on the Different Alignment Algorithms.** To verify that the consistent *rebound* phenomenon occurs in different alignment algorithms in language models, we conduct language models' *rebound* experiments on reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022), direct preference optimization (DPO) (Rafailov et al., 2024), Kahneman-Tversky optimization (KTO) (Ethayarajh et al., 2024), and simple preference optimization (SimPO) (Meng et al., 2024), using an experimental setup similar to the SFT validation described in Section 5.2. Considering that reinforcement learning algorithms cannot be directly applied to pretrained models, our experimental procedure consists of the following two

steps: (a) SFT of the pretrained model using positive sample data of varying scales as *forward alignment*; (b) *inverse alignment* by applying the corresponding reinforcement learning algorithms on the aligned models, using negative sample data of varying scales, where negative samples serve as selected responses and positive samples serve as rejected responses.

We conduct experiments using the Llama2-7B model on the IMDb task (Maas et al., 2011) with four different RL alignment algorithms. The results, shown in Figure 7, indicate that regardless of the alignment algorithm employed, the model exhibits a decline in performance consistent with that observed in SFT as the amount of positive data increases. Moreover, the rate of this performance degradation accelerates with the increasing volume of positive data. The results further confirm that the broad applicability of language models *rebound* across different alignment algorithms.

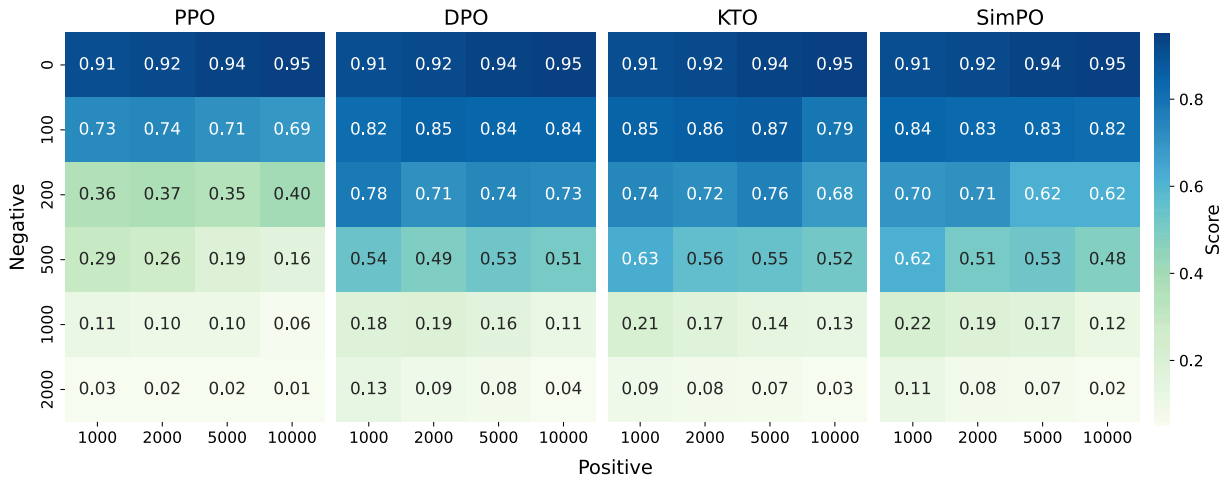


Figure 7: **Experimental results for validating rebound across different alignment algorithms.** The heatmaps from left to right represent the performance of the PPO, DPO, KTO, and SimPO algorithms after fine-tuning with varying amounts of positive and negative data. Darker colors indicate better model performance in the heatmaps. Models trained with more **positive** data perform better initially tend to exhibit worse performance after fine-tuning with **negative** data. The conclusion holds regardless of the alignment algorithm used.

**Ablations on the Different Evaluation Metrics.** In the main experiments discussed in Section 5.2, we primarily adopt specific evaluation scores to evaluate the *rebound* phenomenon. Using the score models corresponding to alignment targets as the evaluation metrics is the most direct and important method for the task evaluation. However, to further reinforce the robustness of our findings and address the generality of the *rebound* phenomenon, we acknowledge the value of incorporating broader distribution-level metrics beyond task-specific scoring models. In particular, we consider the Kullback-Leibler divergence between aligned models and base models as an additional evaluation metric.

Specifically, we conducted the evaluation through the following procedure: a) We first measure the KL divergence between the aligned model and the base model before and after fine-tuning the base model with safety data. b) We then apply fine-tuning with unsafe data until the KL divergence between the models decreased to a sufficiently small value  $\epsilon$  (where  $\epsilon = 1 \times 10^{-2}$ ), and record the amount of unsafe data required as an indicator of the difficulty for the model reverting to its pre-trained distribution. The experimental results are reported in Table 3. Under different models and data scales, we observe that only a small amount of unsafe data is needed for a positively fine-tuned model to revert to the pre-training distribution in terms of KL divergence. Furthermore, the larger the amount of safe fine-tuning, the less unsafe data are required. This observation is consistent with the conclusions drawn in Figure 4.

**Ablations on the Reverse Finetuning Settings.** To rule out the influence of positive data on the *rebound* phenomenon in language models, we conducted a reverse experimental setup: negative data are used during the SFT stage, while positive data are applied during the *inverse alignment* stage. The experimental results, presented in Figure 8, demonstrate that *elasticity* in language models persists under this reverse setting, exhibiting a consistent trend where larger model sizes correspond to greater *elasticity*. The

Base Models	Positive Data Amount			
	1000	2000	5000	10000
Llama2-7B	0.21	0.22	0.26	0.27
Gemma-2B	0.18	0.21	0.24	0.25

(a) KL divergence between models fine-tuned with varying amounts of safety data and the initial models.

Base Models	Positive Data Amount			
	1000	2000	5000	10000
Llama2-7B	961	844	801	729
Gemma-2B	923	853	709	598

(b) Unsafe data amount needed for KL divergence between fine-tuned model and pre-trained model to drop below  $\epsilon$ .

Table 3: **Experimental results for validating rebound on the KL divergence metrics.** The results show that, under different models, the larger the amount of safe fine-tuning, the less unsafe data is required to revert to the pre-training distribution in terms of KL divergence, providing evidence for the existence of *Rebound*.

observation aligns with and further supports the results reported in Figure 5, highlighting the robustness of *rebound* phenomenon across different experimental configurations.




Figure 8: **Reverse Fine-tuning Results on IMDb.** Each subfigure from left to right shows the changes in LLMs with parameter sizes of 0.5B, 4B, and 7B, respectively. Models trained with more **negative** data initially perform worse, but perform better after fine-tuning with **positive** data. As the model size increases, the performance of the aligned model deteriorates more rapidly after fine-tuning with **positive** data.

### B.3 Additional Experiment of Internal Factor of Language Models’ *Rebound*

**Analysis of the Model Size Scale under Different Alignment Algorithms.** To verify whether the *rebound* phenomenon in language models becomes more pronounced as model size increases persists across different alignment algorithms, we conduct experiments using the DPO algorithm. Specifically, we employ the Qwen1.5 model series as the base models and follow an experimental setup similar to that described in Section B.2. We carried out experiments on models with parameters 0.5B, 4B, and 7B. The experimental results, presented in Figure 9, demonstrate that *rebound* remains evident under the DPO algorithm, with its strength increasing as model size grows. This finding is consistent with the results reported in Figure 5 and further highlights the generality of *elasticity* in language models.




Figure 9: **Experimental results for validating rebound increases with model size under DPO fine-tuning.** Each subfigure from left to right shows the changes in LLMs with parameter sizes of 0.5B, 4B, and 7B. As the model size increases, the performance of the aligned model deteriorates more rapidly after fine-tuning with **negative** data.

**Analysis of the Pre-training Data Volume.** We present the experimental results for a broader range of pre-training data volumes in Figure 10. When the pre-training data volume is 0.1T, 0.5T, and 1.0T, the model still demonstrates the phenomenon that *rebound* increases with the volume of pre-training data, which is consistent with the results reported in Figure 6, where the pre-training data volumes range from

2.0T to 3.0T. Moreover, when comparing the experimental results in Figure 6 and Figure 10, we observe a consistent conclusion: as the pre-training data volume increases (from 0.5T to 3.0T), the *rebound* becomes more pronounced. Specifically, when the pre-training data volume increases, the initial performance decline caused by negative data fine-tuning occurs more rapidly, while the subsequent decline slows down. This indicates that larger pre-training data volumes amplify the *rebound* effect in language models.




Figure 10: **Experimental results for validating *rebound* increases with model pre-training data volume.** Each subfigure from left to right shows the changes in pre-training data volumes of 0.1T, 0.5T, and 1.0T. As pre-training data volume increases, aligned model performance deteriorates more rapidly after fine-tuning with **negative** data.

## C Further Discussions of Language Models’ Elasticity

### C.1 Transfer Learning and Language Models’ Elasticity

Transfer Learning is formally defined as follows: consider a domain  $\mathcal{D} = \{\mathcal{X}, P(X)\}$ , which consists of a feature space  $\mathcal{X}$  and a marginal probability distribution  $P(X)$ , where  $X = \{x_1, \dots, x_n\} \in \mathcal{X}$ . A task  $\mathcal{T}$  within the domain  $\mathcal{D}$  is defined by a label space  $\mathcal{Y}$  and a predictive function  $f(\cdot)$ . Given a source domain  $\mathcal{D}_S$  and its corresponding task  $\mathcal{T}_S$ , as well as a target domain  $\mathcal{D}_R$  and its task  $\mathcal{T}_R$ , transfer learning aims to improve the target predictive function  $f_R(\cdot)$  by leveraging knowledge from  $\mathcal{D}_S$  and  $\mathcal{T}_S$ , where either  $\mathcal{D}_S \neq \mathcal{D}_R$  or  $\mathcal{T}_S \neq \mathcal{T}_R$  (Zhuang et al., 2020; Weiss et al., 2016). From the perspective of transfer learning, the pre-training and fine-tuning process of language models can be described as follows: Given labeled training data from a source distribution  $\mathcal{D}^s = \{(x_n, y_n) \sim p\}_{n=1}^{N_s}$  and data from a target distribution  $\mathcal{D}^t = \{(x_n, y_n) \sim q\}_{n=1}^{N_t}$ , we first minimize the empirical risk on the source distribution to fit a model:  $f^s = \operatorname{argmin}_f \hat{R}(f, \mathcal{D}^s)$ , where  $\hat{R}(f, \mathcal{D}^s) = \mathbb{E}_{(x,y) \sim p}[l(y, f(x))]$  represents the empirical risk of model  $f$  on the source distribution  $p$ . Subsequently, the model is adapted to the target distribution by minimizing,  $f^t = \operatorname{argmin}_f \hat{R}(f, \mathcal{D}^t) + \lambda \|f - f^s\|$ , where  $\|f - f^s\|$  measures some distance between the two functions, and  $\lambda \geq 0$  is a regularization parameter (Murphy, 2023).

Specifically in the context of language models, transfer learning can almost encompass all models trained under the pre-training-finetuning paradigm. However, transfer learning does not capture the specific details of the model training process, nor does it explicitly incorporate the established theories that explain the varying degrees of difficulty in task transfer for LLMs. Therefore, while transfer learning provides a perspective to analyze various phenomena in LLMs from the standpoint of traditional machine learning, it does not offer concrete tools to explain these *elasticity* phenomena.

Furthermore, in traditional transfer learning settings, positive and negative samples can only be considered as different partitions of the same dataset in classification tasks. However, in our case, these samples are used for autoregressive generation training, where the generation on positive samples and negative samples follows mutually conflicting input-output distributions. This is fundamentally different from classification tasks. Specifically for safe-unsafe generation, steering a language model’s generation from an unsafe input-output distribution to the opposite safe input-output distribution is a major challenge in generative language model research. This cannot be adequately explained by merely considering different partitions of the same data distribution.

These considerations lead us to conclude that the observed phenomenon in our work represents a novel behavior that cannot be fully accounted for by existing transfer learning theories. It warrants finer-grained theoretical and experimental investigation. Consequently, we propose the concept of *elasticity* to formally name this new phenomenon observed in language models and introduce corresponding theoretical

frameworks and empirical methodologies to examine it. This perspective is empirically supported by the experimental results shown in Table 1. Taking the Alpaca task as an example, we examine the process in which  $\theta_1$  and  $\theta_2$  are aligned with  $\theta_3$ . Since  $\theta_2$  has acquired more knowledge about the SFT distribution compared to  $\theta_1$ , it should, from the perspective of transfer learning, find it easier to learn the distribution of  $\theta_3$  than  $\theta_1$ . However, as shown in the experimental results above, this is not supported, suggesting that transfer learning alone cannot fully account for the observed *resistance* phenomenon.<sup>2</sup>

## C.2 Discussions of Practical Steps to Mitigate *Inverse Alignment* Risks

The emphasis on *elasticity* phenomenon in language model alignment highlights a crucial challenge for ensuring the safety and robustness of open-source models throughout their entire lifecycle. As our study demonstrates, the resistance of language models to alignment adjustments is fundamentally rooted in the substantial disparity in data volumes across different training phases. This insight suggests a practical mitigation strategy: customizing the scale of synthetic data through the elasticity mechanism offers a feasible pathway for the development of more robust alignment algorithms in the future.

- We discover that the resistance of language models to alignment is essentially due to the significant differences in data volume across various training processes. Therefore, a straightforward idea is to ensure that the training data volumes corresponding to different alignment targets are as similar as possible during the alignment process. This approach helps avoid resistance effects on alignment targets with smaller training data volumes due to subsequent perturbations.
- The elasticity theorem provides a feasible method for customizing data ratios, thereby quantitatively measuring the amount of various types of data needed for a model to meet expected goals. Specifically, for an aligned language model, during subsequent alignment processes, the elasticity theorem indicates that there is an inherent loss of elasticity dependent on data volume ratios when learning unrelated distributions. Therefore, to ensure that previously aligned targets remain satisfied during subsequent alignment, the training objective can be transformed into a constrained optimization problem to quantitatively calculate the data volume required for new target alignment.
- The current challenge in implementing algorithms based on the elasticity theory, mainly lies in the lack of a sufficient theoretical foundation to characterize the specific features of current datasets. This makes it difficult to distinguish the independent and differently distributed premises of different datasets in the elasticity theory. Since real-world datasets often contain fused features, a possible future research direction is to finely characterize the features of datasets with fused characteristics.

## C.3 Discussions of Quantitatively Characterize the Impact of Dataset Size on *Elasticity*

Considering that, as indicated in Theorem 4.2, the difference in dataset size is a key factor contributing to the *elasticity* phenomenon, quantitatively characterizing the impact of dataset size on *elasticity* is indeed crucial for both theoretical understanding and practical alignment design. In this section, we present preliminary explorations, based on experimental observations, on quantifying the influence of dataset size on elasticity, as well as the current limitations in achieving this goal.

By synthesizing observations from Figure 6 in the main text and Figure 10 in the appendix, we find that for the tinyllama series models, when the pre-training data volume is merely 0.1T, there is almost no significant resistance phenomenon observed; the positive level of the model's output does not change significantly with an increase in negative data. However, when the pre-training data volume reaches 0.5T, the resistance phenomenon of the language model becomes notably observable. Considering the continuity between pre-training data and model resistance, this suggests that the significant change point of the resistance phenomenon should lie between 0.1T and 0.5T.

Regrettably, there are still challenges in training a series of models with continuously varying pre-training data volumes, and we currently lack sufficiently fine-grained training data samples to conduct such precise series model training. This limitation prevents us from accurately determining the critical

---

<sup>2</sup>As Reviewer qznt's suggestion, we further discuss the relationship between transfer learning and LLMs' elasticity in this section.

point where the resistance phenomenon becomes significant. Therefore, although we can suggest a possible range, accurately quantifying this threshold in practice remains a considerable challenge. We hope that future research will further address this issue by developing more controllable model series and fine-grained pre-training datasets to enable a clearer, quantitative characterization of *elasticity*.