

Resonant Context Anchoring: Decoupling Attention Routing and Signal Gain at Inference Time

Mingkuan Zhao^{1*}, Yide Gao^{1*}, Wentao Hu¹, Suquan Chen¹, Tianchen Huang^{2†},
Zhenhua An¹, Zetao Chang³, Xiayu Sun¹, Yuheng Min⁴

¹Xi'an Jiaotong University, ²University of Science and Technology of China,
³Tongji University, ⁴Tsinghua University

{mingkuanzhao, yidegao, wentao_hu, suquanchen, zhenhuan, xiayu_sun}@stu.xjtu.edu.cn,
tchuang@mail.ustc.edu.cn, 2534034@tongji.edu.cn, minyh24@mails.tsinghua.edu.cn

Abstract

Large Language Models (LLMs) frequently exhibit “contextual disregard” when faced with input evidence that conflicts with their internal parametric memory, leading to persistent factual hallucinations. Existing mitigation strategies primarily rely on suppressing specific neuron activations or employing computationally expensive contrastive decoding mechanisms, which often result in increased perplexity or significantly elevated inference latency. To address these limitations, we propose **Resonant Context Anchoring (RCA)**, a lightweight inference-time intervention method grounded in the perspective of residual stream signal dynamics. RCA aims to resolve the signal attenuation of external evidence during its propagation through deep networks. The core mechanism involves the orthogonal decoupling of routing logic and information magnitude within the self-attention module. By utilizing raw pre-softmax attention scores as an instantaneous metric of semantic alignment, we construct a dynamic gain field via non-linear rectification to selectively amplify the norms of value vectors corresponding to context tokens, without altering the attention probability distribution. This mechanism effectively elevates the signal-to-noise ratio (SNR) of input evidence within the residual stream mixture, thereby robustly anchoring the generation trajectory to the truthful context during inference. Extensive experiments on the Llama-3 model series demonstrate that RCA significantly improves contextual faithfulness across multiple factual consistency and strong knowledge-conflict tasks, effectively suppressing parametric hallucinations. Furthermore, results confirm that as a training-free and computationally negligible plug-and-play module, RCA achieves a Pareto improvement in faithfulness and fluency while maintaining the model’s general language understanding capabilities. Our code is available at [https://](https://github.com/yidGao/RCA-Implementation)

github.com/yidGao/RCA-Implementation.

1 Introduction

Despite the remarkable generalization capabilities demonstrated by Large Language Models (LLMs) (Brown et al., 2020; OpenAI et al., 2024; He et al., 2024; Li et al., 2024c) in natural language generation tasks, their reliance on parametric memory—internalized during pre-training—renders them susceptible to factual errors, particularly when processing input evidence that contradicts internal priors. This phenomenon, termed parametric hallucination (Ji et al., 2023), reveals a fundamental tension between a model’s internalized beliefs and the provided context. Recent scholarship indicates a persistent robustness failure: even when presented with accurate and relevant evidence, LLMs frequently exhibit “contextual disregard,” prioritizing intrinsic parametric priors over external information. This tendency to override input facts with internal biases significantly undermines the reliability of LLMs in knowledge-intensive and high-stakes reasoning scenarios.

Current approaches to mitigating these hallucinations predominantly fall into two categories: supervised fine-tuning and inference-time intervention. While Supervised Fine-Tuning (SFT) can improve instruction following through targeted data alignment, it incurs high computational costs and risks overfitting to specific conflict patterns, thereby limiting generalizability. Conversely, inference-time interventions attempt to modulate generation behavior without updating model weights. Contrastive Decoding (CD) strategies, for instance, rectify output distributions by contrasting logits generated with and without context. However, these methods typically necessitate multiple forward passes, doubling inference latency and rendering them impractical for real-time applications. Alternative methods based on Activation Engineering seek to identify and suppress specific attention

*These authors contributed equally to this work.

†Corresponding author.

heads or neurons responsible for retrieving parametric memory. Yet, such coarse-grained suppression often disrupts the underlying syntactic structure and semantic coherence, resulting in increased perplexity and degraded linguistic fluency.

In this work, we revisit the phenomenon of contextual disregard through the lens of signal dynamics within the residual stream. We hypothesize that the model’s failure in conflict scenarios stems not from an inability to identify relevant context (a routing error), but from the insufficient energy of the contextual signal during propagation through deep networks (a gain deficit). Within the Transformer architecture (Vaswani et al., 2017), the residual stream can be modeled as a linear superposition of parametric and contextual subspaces. Even if the query vector correctly attends to context tokens, if the norms of the corresponding Value vectors are significantly smaller than the parametric noise vectors injected by the Multi-Layer Perceptrons (MLPs) (Geva et al., 2021), the effective information becomes submerged, leading to a critical drop in the Signal-to-Noise Ratio (SNR). Consequently, the solution to factual anchoring lies not in the destructive suppression of internal memory, but in the selective amplification of the contextual signal magnitude. This conceptual shift in geometric projection is illustrated in Figure 1.

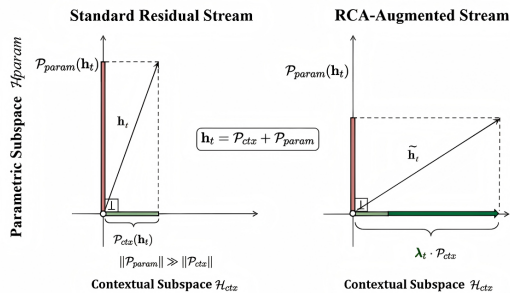


Figure 1: **Geometric interpretation of subspace projection.** Left: In the standard residual stream, the parametric projection \mathcal{P}_{param} dominates the energy, causing the state \mathbf{h}_t to align with internal priors. Right: RCA applies a gain λ_t to the contextual projection, effectively rotating the resultant vector $\tilde{\mathbf{h}}_t$ towards the contextual manifold without altering the underlying basis.

To this end, we propose **Resonant Context Anchoring (RCA)**, a lightweight inference-time intervention method. The core contribution of RCA is the orthogonal decoupling of “routing decisions” and “information gain” within the self-attention mechanism. Specifically, we utilize raw

pre-softmax attention scores as an instantaneous metric of semantic alignment to drive a non-linear rectification mechanism. During inference, when a high semantic correlation (resonance) between the query and context tokens is detected, RCA dynamically amplifies the norms of the corresponding Value vectors. This process significantly enhances the weight of input evidence within the residual stream without altering the attention probability distribution, thereby preserving both syntactic correctness and factual faithfulness.

We conduct extensive evaluations on the Llama-3 model series. Empirical results demonstrate that RCA significantly improves contextual faithfulness in summarization and strong knowledge-conflict benchmarks (e.g., NQ-Swap), effectively mitigating hallucinations driven by prior bias. Crucially, as a training-free plug-and-play module introducing only negligible element-wise operations, RCA achieves a Pareto improvement in faithfulness and fluency without incurring the latency penalties associated with iterative decoding methods.

2 Related Work

Contextual grounding has emerged as a fundamental framework for mitigating factual hallucinations in Large Language Models (LLMs) by anchoring outputs in external evidence (Lewis et al., 2021). However, a persistent challenge arises when input evidence contradicts the model’s internalized parametric memory, leading to “contextual disregard” where the model prioritizes intrinsic priors over provided facts (Longpre et al., 2021). Previous scholarship has highlighted the limitations of LLMs in handling complex conflicting evidence and long-tail knowledge (Chen et al., 2022; Mallen et al., 2023), necessitating adaptive augmentation strategies that dynamically calibrate the reliance on external context based on generation confidence (Wang et al., 2025a). These findings underscore the need for decoding mechanisms that robustly anchor generation to the input context during knowledge conflicts.

The foundation of the hallucination problem lies in the pre-training paradigm of modern LLMs. Models such as the TeleChat series (Wang et al., 2024; He et al., 2024; Wang et al., 2025b; Liu et al., 2025) and the Tele-FLM family (Li et al., 2024c,b) demonstrate that bilingual and large-scale pre-training over trillions of tokens endows models with strong generalization, yet simultaneously en-

trenches parametric priors that may conflict with externally provided evidence at inference time. These capable base models have been deployed across a broad spectrum of knowledge-intensive downstream tasks: TableReasoner (Xiong et al., 2025) advances structured table reasoning by leveraging LLMs’ semantic understanding, MR-UIE (Li et al., 2025) applies multi-perspective reinforcement learning to universal information extraction, and Zhao et al. (2025a) enhance mathematical reasoning through computation logic graphs. The LLMSR pipeline (Xing et al., 2025) further addresses the challenge of structured reasoning data construction, combining gradient-reward policy optimization with iterative self-refinement. Collectively, these works underscore that while LLMs can reliably internalize and apply broad world knowledge, their tendency to override contextual evidence with parametric memory represents a critical barrier to deployment in high-reliability, knowledge-conflict scenarios.

Complementary to faithfulness research, recent work has revisited the internal routing logic of the attention mechanism from both efficiency and structural perspectives. Zhao et al. (2025b) propose a sparse attention strategy that achieves competitive performance without sacrificing inference speed, demonstrating that not all attention heads contribute equally to the information routing process. In parallel, Mosaic Pruning (Hu et al., 2025) introduces a hierarchical pruning framework for Mixture-of-Experts (MoE) models that generalizes across model families, revealing that structured sparsity in attention and feed-forward blocks can be exploited for efficient deployment without significant performance loss. These findings are closely aligned with the signal-level perspective adopted in our work: if attention routing can be decoupled from signal magnitude for the purpose of efficiency, the same decoupling principle can be leveraged to selectively amplify contextual evidence signals during inference—the central insight behind RCA.

Inference-time intervention strategies represent a primary category of techniques for enhancing model faithfulness. Contrastive Decoding (CD) optimizes reasoning quality by calibrating output distributions across different model configurations (O’Brien and Lewis, 2023). Notably, Context-aware Decoding (CAD) enforces context adherence by amplifying the logit difference between conditional and unconditional generations (Shi et al., 2024). Other methods, such as DoLa, contrast ma-

ture and premature layer representations to improve factuality (Chuang et al., 2024), while Inference-Time Intervention (ITI) injects linear directions associated with truthfulness into specific attention heads (Li et al., 2024a). Although effective, these approaches often introduce significant computational overhead or require the training of auxiliary probes.

Advances in activation engineering and mechanistic interpretability offer a physical basis for precisely steering model behavior (Turner et al., 2024). Research has demonstrated that latent steering vectors can be extracted to manipulate the semantic space of pretrained decoders without fine-tuning (Subramani et al., 2022). Techniques like Contrastive Activation Addition (CAA) apply displacement vectors in the residual stream to guide specific behaviors such as factuality (Rimsky et al., 2024). Structurally, feed-forward layers are identified as key-value memories for parametric facts (Geva et al., 2021), while induction heads facilitate the replication of contextual information across layers (Olsson et al., 2022). Recent mathematical perspectives on Transformer dynamics (Geshkovski et al., 2025) and methods for editing implicit assumptions (Orgad et al., 2023) further support the feasibility of fine-grained signal modulation within the residual stream.

The maturation of factuality evaluation frameworks provides quantitative metrics for assessing intervention efficacy. FACTScore enables fine-grained evaluation through atomic fact decomposition (Min et al., 2023), while AlignScore utilizes a unified alignment function to measure factual consistency across diverse tasks (Zha et al., 2023). These metrics, combined with systematic analyses of LLMs as factual reasoners (Laban et al., 2023), form the experimental foundation for modern research. Building upon these foundations, our proposed RCA method introduces a lightweight intervention that recalibrates the signal-to-noise ratio in the residual stream, offering an efficient alternative to contrastive decoding while preserving architectural integrity better than coarse-grained activation suppression.

3 Methodology

In this section, we formally define the algorithmic framework of **Resonant Context Anchoring (RCA)**. We reconstruct the self-attention mechanism of the Transformer architecture as a signal

transmission system equipped with adaptive spectral gain control. Based on the theory of vector space decomposition within the residual stream, we derive a non-linear rectification mechanism driven by semantic correlation. This approach aims to dynamically calibrate the energy distribution of the contextual subspace by orthogonally decoupling the routing decision logic from the signal intensity.

3.1 Subspace Decomposition and SNR Dynamics in Residual Streams

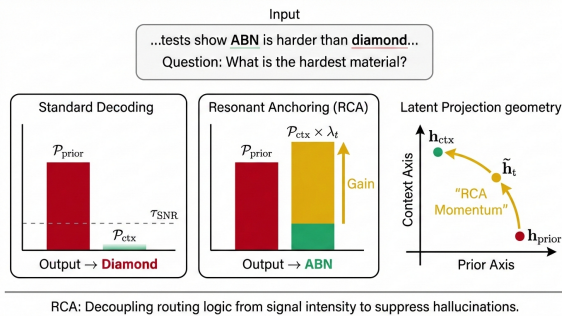


Figure 2: **Mechanistic illustration of signal dynamics.** Left: Standard decoding fails when contextual evidence is submerged by priors. Middle: RCA restores the signal-to-noise ratio. Right: The resulting "RCA Momentum" in latent space.

At the l -th layer of the Transformer architecture, the hidden state vector $\mathbf{h}_t^{(l)} \in \mathbb{R}^d$ can be formalized as a linear superposition of distinct information sources (Elhage et al., 2021) within a high-dimensional latent space. For the t -th time step in an autoregressive generation process, the state update follows the residual connection rule, driven jointly by the Multi-Head Attention (MHA) and Feed-Forward Network (FFN) sub-layers. We decompose this state space \mathbb{R}^d into two competing subspaces: the Contextual Subspace $\mathcal{H}_{ctx} \subset \mathbb{R}^d$, spanned by external evidence within the current context window, and the Parametric Subspace $\mathcal{H}_{param} \subset \mathbb{R}^d$, implicitly encoded by the model’s pre-trained weight matrices. Consequently, the residual vector \mathbf{h}_t can be expressed as the superposition of two orthogonal projection components:

$$\mathbf{h}_t \approx \mathcal{P}_{ctx}(\mathbf{h}_t) + \mathcal{P}_{param}(\mathbf{h}_t) \quad (1)$$

where \mathcal{P}_{ctx} and \mathcal{P}_{param} denote the projection operators onto the contextual and parametric subspaces, respectively. The mechanism underlying hallucinations or knowledge conflicts can be interpreted from a signal processing perspective as a critical imbalance in the Signal-to-Noise Ratio (SNR) within

the residual stream, leading to what we term “signal submergence” (Figure 2). Specifically, when processing a specific query \mathbf{q}_t , if the L_2 norm of the parametric component significantly dominates the residual stream, satisfying the inequality condition:

$$\|\mathcal{P}_{param}(\mathbf{h}_t)\|_2 \gg \|\mathcal{P}_{ctx}(\mathbf{h}_t)\|_2 \quad (2)$$

the model’s generation trajectory will collapse onto the internal prior manifold, leading to the disregard of external evidence. The standard attention mechanism employs the Softmax function to achieve probabilistic normalization; however, this operation only constraints the routing distribution of information aggregation without explicitly controlling the scalar magnitude of the aggregated information. Even if the attention weights correctly point to context tokens, if the raw magnitude of the value vector \mathbf{v}_i is insufficient to counteract the high-norm parametric noise injected by the FFN layers in deep networks, the effective signal will attenuate during hierarchical propagation. Therefore, our objective is to construct a gain operator $\mathcal{G}(\cdot)$ that selectively enhances the modified contextual projection norm $\|\mathcal{P}_{ctx}(\mathcal{G}(\mathbf{h}_t))\|$ while maintaining the invariant routing distribution $\alpha_t \in \Delta^{L-1}$.

3.2 Semantic Spectrum Analysis and In-Situ Value Modulation

The detailed internal architecture of the RCA module, which implements the decoupling of routing logic and signal gain, is depicted in Figure 3.

To achieve the aforementioned gain objective, RCA introduces a non-linear modulation mechanism based on the semantic alignment spectrum. This mechanism operates directly on the value projection stage of the attention layer, realizing the orthogonal decoupling of routing logic and signal gain. First, for the h -th attention head, we compute the raw correlation score between the current query vector $\mathbf{q}_t \in \mathbb{R}^{d_k}$ and the context key vector matrix $\mathbf{K} \in \mathbb{R}^{L \times d_k}$. This unnormalized dot product constitutes the raw spectrum distribution of semantic alignment $\mathbf{s}_t \in \mathbb{R}^L$, where the i -th component is defined as:

$$s_{t,i} = \frac{\mathbf{q}_t^\top \mathbf{k}_i}{\sqrt{d_k}} \quad (3)$$

This scalar precisely quantifies the degree of collinearity between the query intent and the i -th context token in the semantic space. Unlike the standard attention mechanism which utilizes \mathbf{s}_t solely for computing Softmax probabilities, RCA

Resonant Context Anchoring (RCA)

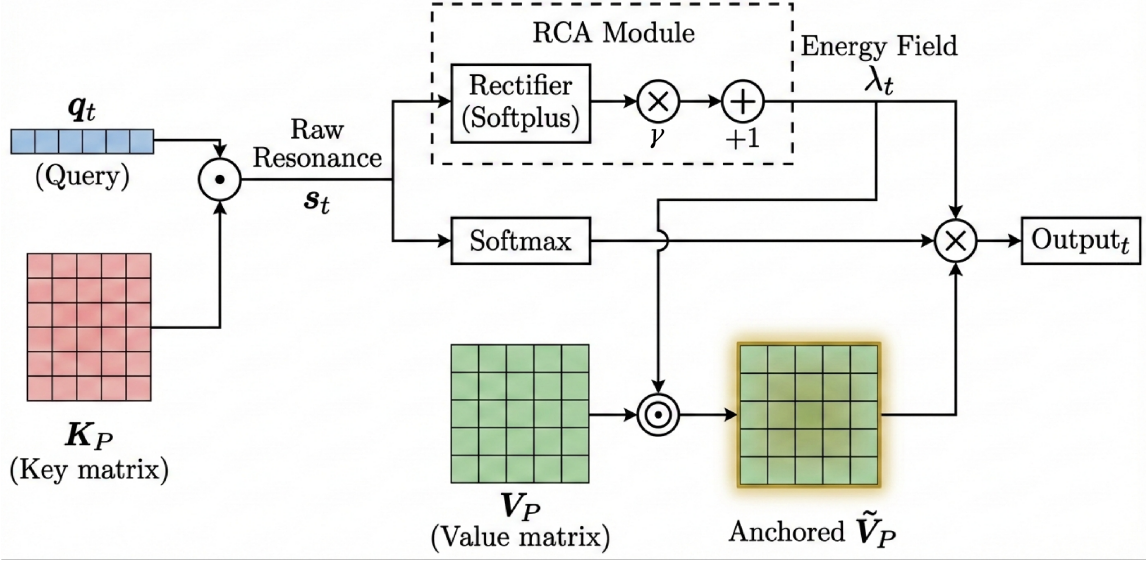


Figure 3: **The architecture of Resonant Context Anchoring (RCA).** The module decouples routing logic (Softmax) and information magnitude. It computes an energy field λ_t via a non-linear rectifier (Softplus) applied to the raw resonance s_t , which is then used to modulate the value matrix V_P to produce the anchored representation \tilde{V}_P .

introduces a parallel rectification path to construct a dynamic gain field. We define a non-linear rectification function $\psi : \mathbb{R} \rightarrow \mathbb{R}^+$ to extract the "resonance intensity." To ensure the smoothness of the manifold and avoid gradient truncation, we employ the Softplus operator with a sensitivity coefficient γ , constructing a scalar gain $\lambda_{t,i}$ for each token:

$$\lambda_{t,i} = 1 + \gamma \cdot \psi(s_{t,i}) = 1 + \gamma \cdot \ln(1 + \exp(s_{t,i})) \quad (4)$$

The gain function $\lambda_{t,i}$ possesses rigorous mathematical properties: as $s_{t,i} \rightarrow -\infty$ (indicating semantic orthogonality or negative correlation), $\lambda_{t,i} \rightarrow 1$, causing the system to degenerate into an identity mapping and ensuring zero interference with noise; when $s_{t,i} > 0$ (indicating semantic resonance), $\lambda_{t,i}$ exhibits quasi-linear monotonic growth with respect to correlation. Subsequently, we apply this gain field to the value vector matrix $\mathbf{V} \in \mathbb{R}^{L \times d_v}$ via an in-situ Hadamard product modulation. The modified value vector $\tilde{\mathbf{v}}_i$ is defined as:

$$\tilde{\mathbf{v}}_i = \lambda_{t,i} \cdot \mathbf{v}_i \quad (5)$$

Finally, the output of the attention head \mathbf{o}_t is computed via the standard probability-weighted summation, but taking the energy-enhanced value vectors as input. The complete computational form

can be expanded as:

$$\mathbf{o}_t = \sum_{i=1}^L \text{Softmax}(s_t)_i \cdot [(1 + \gamma \cdot \text{Softplus}(s_{t,i})) \cdot \mathbf{v}_i] \quad (6)$$

From the perspective of vector synthesis, this operation is equivalent to applying a momentum term controlled jointly by γ and $s_{t,i}$ in the direction of the context-relevant gradient, forcing the residual stream to converge towards the ground-truth evidence manifold.

3.3 Asymptotic Analysis of Local SNR Gain

To theoretically validate the effectiveness of RCA, we perform an asymptotic analysis on the gain boundary of the local Signal-to-Noise Ratio (SNR). Let the effective signal energy E_{sig} at step t be the weighted norm of value vectors within the relevant context set \mathcal{I}_{rel} , and the noise energy E_{noise} be the norm of irrelevant context and parametric bias. Under the standard attention mechanism, the base SNR is given by $\text{SNR}_{base} = E_{sig}/E_{noise}$. Upon introducing RCA, due to the non-negativity and monotonicity of the rectification function ψ , for any $i \in \mathcal{I}_{rel}$, the high alignment with the query vector implies $s_{t,i} > 0$, thereby ensuring $\lambda_{t,i} > 1$. The modified signal energy \hat{E}_{sig} can be expressed

as:

$$\tilde{E}_{sig} = \left\| \sum_{i \in \mathcal{I}_{rel}} \alpha_{t,i} (1 + \gamma \bar{\psi}) \mathbf{v}_i \right\| \quad (7)$$

Utilizing the triangle inequality and its reverse property, and assuming the resonance intensity distribution within the relevant context is concentrated, we have $\tilde{E}_{sig} \approx (1 + \gamma \bar{\psi}) E_{sig}$, where $\bar{\psi}$ denotes the average resonance intensity of relevant tokens. Meanwhile, since RCA does not intervene in the FFN layer outputs, and for irrelevant tokens $j \notin \mathcal{I}_{rel}$, $s_{t,j}$ is typically negative leading to $\lambda_{t,j} \approx 1$, the noise energy term remains quasi-constant, i.e., $\tilde{E}_{noise} \approx E_{noise}$. Consequently, the post-intervention SNR satisfies the following inequality:

$$\begin{aligned} \text{SNR}_{RCA} &\approx \frac{(1 + \gamma \bar{\psi}) E_{sig}}{E_{noise}} \\ &= (1 + \gamma \bar{\psi}) \text{SNR}_{base} > \text{SNR}_{base} \end{aligned} \quad (8)$$

This inequality mathematically guarantees that RCA monotonically enhances the model’s sensitivity to external evidence. By adjusting the hyperparameter γ , we can explicitly control the gain magnitude of the SNR, thereby effectively escaping the hallucination attractor without disrupting the model’s language modeling distribution.

4 Experiments

In this section, we present a comprehensive empirical evaluation to assess the effectiveness, robustness, and parametric sensitivity of Resonant Context Anchoring (RCA). Our experimental design aims to systematically address three core research questions: First, can the proposed method significantly enhance contextual faithfulness in summarization tasks without compromising generation quality? Second, in adversarial scenarios characterized by strong conflicts between parametric memory and external evidence, does the method effectively suppress hallucinations and enforce context adherence? Third, does the method exhibit general safety, maintaining the model’s fundamental capabilities in tasks lacking relevant context? All experiments are conducted on the Llama-3-8B-Instruct and Llama-3-70B-Instruct models (Grattafiori et al., 2024), benchmarking RCA against standard greedy decoding strategies.

4.1 Experimental Setup and Benchmarks

To rigorously evaluate model performance across varying cognitive loads, we constructed a diverse suite of benchmarks encompassing single-document summarization, knowledge-conflict QA, and general commonsense QA. For the assessment of contextual faithfulness, we utilized the XSum dataset (Narayan et al., 2018), employing FactKB (Feng et al., 2023) and AlignScore (Zha et al., 2023) as primary metrics for factual consistency between the generated text and source documents, complemented by ROUGE-L (Lin, 2004) to evaluate content coverage. In the more challenging domain of knowledge conflict, we selected the NQ-Swap (Longpre et al., 2021) and MemoTrap (McKenzie et al., 2024) datasets. NQ-Swap constructs counter-factual contexts by substituting key entities (e.g., swapping countries or capitals) within Wikipedia passages, forcing the model to choose between internal correct knowledge and external altered context; we report Exact Match (EM) as the core metric for context adherence. MemoTrap tests the model’s ability to resist rote memorization when completing counter-intuitive proverbs. For general capability evaluation, we employed the closed-book settings of TriviaQA (Joshi et al., 2017) and TruthfulQA (Lin et al., 2022) to verify that the inference-time intervention does not induce catastrophic forgetting of general world knowledge. Implementation-wise, RCA is integrated as a parameter-free module within the Hugging Face Transformers library (Wolf et al., 2020), with inference performed on NVIDIA A100 GPU clusters.

4.2 Contextual Faithfulness and Conflict Resolution

The main experimental results, presented in Table 1, indicate that RCA achieves performance significantly superior to the baseline in enhancing contextual faithfulness. In the XSum summarization task, the Llama-3-8B model with RCA shows an increase in the FactKB score from a baseline of 47.61 to 50.14, alongside a concurrent rise in AlignScore from 58.20 to 59.45. Crucially, this improvement in factual consistency is not attained at the expense of content richness; the ROUGE-L score remains stable or improves slightly (increasing from 19.90 to 20.03). Similar trends are observed in the larger Llama-3-70B model, where FactKB improves from 61.32 to 61.72 and AlignScore from

Table 1: **Main Results on Faithfulness and Knowledge Conflict.** The table reports the performance of standard greedy decoding (Baseline) versus RCA across two model scales. RCA demonstrates consistent improvements across all metrics. Notably, it achieves a Pareto improvement by enhancing faithfulness (FactKB, AlignScore) without degrading content coverage (ROUGE-L). (\uparrow indicates higher is better).

Task	Metric	Llama-3-8B		Llama-3-70B	
		Baseline	RCA	Baseline	RCA
Contextual Faithfulness (XSum)					
Summarization	FactKB (\uparrow)	47.61	50.14	61.32	61.72
	AlignScore (\uparrow)	58.20	59.45	65.10	65.88
	ROUGE-L (\uparrow)	19.90	20.03	22.41	22.21
Knowledge Conflict Resolution					
NQ-Swap	Exact Match (\uparrow)	60.62	64.54	76.11	77.46
MemoTrap	Micro Acc (\uparrow)	64.40	65.77	66.52	77.35
	Macro Acc (\uparrow)	65.86	66.69	68.47	75.58

65.10 to 65.88. These findings corroborate that RCA does not merely circumvent hallucinations by truncating outputs, but rather successfully steers the model to attend more intensively to key information fragments within the source document, thereby realizing a Pareto improvement between faithfulness and informativeness.

In tasks characterized by severe conflicts between parametric memory and external evidence, RCA exhibits a robust capability to suppress internal priors. On the NQ-Swap dataset, the Llama-3-8B model under standard decoding achieves an Exact Match rate of only 60.62%, implying that in approximately 40% of cases, the model prioritizes internal correct facts over the specified context. The introduction of RCA boosts the Exact Match rate to 64.54%, an absolute improvement of 3.92%. This trajectory is maintained in the Llama-3-70B model, where the Exact Match rate rises from 76.11% to 77.46%. Results from the MemoTrap dataset further substantiate this conclusion: for Llama-3-70B, the Micro Accuracy surges from 66.52% to 77.35%, and Macro Accuracy improves from 68.47% to 75.58%. These data provide compelling evidence that by dynamically amplifying contextual signals resonating with the query in the residual stream, RCA effectively breaks the hold of parametric attractors, enabling the model to anchor robustly to external evidence even amidst intense conflict.

Table 2: **Sensitivity Analysis of γ on NQ-Swap.** The Exact Match (EM) scores illustrate the non-linear dynamics. Performance peaks in the [0.04, 0.05] interval ("sweet spot") and degrades significantly beyond 0.08.

γ Setting	Llama-3-8B	Llama-3-70B
Baseline ($\gamma = 0$)	60.62	76.11
$\gamma = 0.02$	62.15	76.85
$\gamma = 0.04$	64.54	77.20
$\gamma = 0.05$	64.10	77.46
$\gamma = 0.08$	62.80	76.90
$\gamma = 0.10$	58.40	73.50
$\gamma = 0.12$	52.10	68.20

4.3 Parameter Sensitivity and Non-linear Dynamics

The resonance sensitivity coefficient γ serves as the core hyperparameter controlling the intensity of the RCA intervention. We conducted a granular grid search and ablation analysis (detailed in Table 2) to investigate its non-linear impact on model behavior. The experiments reveal a clear and consistent performance evolution trajectory. In the interval where γ ranges from 0 to 0.04, all key metrics exhibit a monotonic upward trend as the coefficient increases. For instance, with Llama-3-8B on NQ-Swap, increasing γ from 0.02 to 0.04 raises the Exact Match rate from 62.15% to a peak of 64.54%. The model performance reaches a global optimum plateau when γ lies between 0.04 and 0.05. Specifically, Llama-3-8B performs best at $\gamma = 0.04$, while Llama-3-70B peaks at $\gamma = 0.05$ (77.46%), a difference likely attributable to the larger model's greater

embedding space norm, which necessitates slightly stronger external excitation to overcome internal priors.

However, precise physical boundaries exist for parameter settings. When γ exceeds 0.05 and approaches 0.08, performance on select datasets begins to decline; for example, the Exact Match rate for Llama-3-8B regresses to 62.80% at $\gamma = 0.08$. As γ further increases beyond 0.1, the perplexity of the generated text rises sharply, leading to severe degradation in output quality. In the extreme setting of $\gamma = 0.12$, performance drops to 52.10% for Llama-3-8B and 68.20% for Llama-3-70B. Based on this empirical analysis, we fix the optimal configurations for Llama-3-8B and Llama-3-70B at 0.04 and 0.05, respectively, and establish 0.08 as the hard upper threshold for the safe application of this method. This result unveils a distinct "stable region" for the resonance mechanism, within which contextual control can be maximized without introducing adverse side effects.

4.4 Safety and Harmlessness Evaluation

Table 3: **Safety Evaluation on Closed-Book Tasks.** We report performance on tasks with irrelevant or no context provided. The negligible deviation between Baseline and RCA confirms the method’s harmlessness and robust general capabilities.

Task	Metric	Baseline	RCA
TruthfulQA	MC1	38.92	38.92
	MC2	55.64	55.61
TriviaQA	EM	56.58	56.52
PopQA	Acc	26.64	26.59

To verify the harmlessness of RCA, we conducted general QA testing in settings where either no relevant context or irrelevant context was provided (results in Table 3). On the TruthfulQA MC1 (Multiple Choice-Single) metric, the model with RCA scores identically to the baseline (38.92%), and on the MC2 (Multiple Choice-Multiple) metric, the difference is a negligible 0.03% (55.64% vs. 55.61%). Similarly, on the TriviaQA dataset, the Exact Match rate fluctuates only marginally from 56.58% to 56.52%, well within the margin of statistical error. PopQA results exhibit the same trend, with accuracy shifting nominally from 26.64% to 26.59%. The theoretical basis for this result lies in RCA’s resonance detection mechanism: in general QA scenarios, the dot product scores between

the query vector and irrelevant or empty contexts naturally fall into low or negative ranges. After processing through the Softplus rectification, the amplification coefficient approaches 1, effectively rendering the intervention mechanism "dormant." This adaptive activation characteristic ensures that RCA can function as an always-on decoding strategy, enhancing performance on specific context-intensive tasks without necessitating concerns regarding negative impacts on the model’s general linguistic capabilities.

5 Limitations

While Resonant Context Anchoring (RCA) demonstrates significant efficacy in suppressing parametric hallucinations and reinforcing adherence to input evidence, the introduction of the resonance sensitivity coefficient γ necessitates calibration across different model scales or task environments. *However*, empirical evidence indicates that this parameter exhibits cross-model stability, with optimal values remaining largely consistent between Llama-3-8B and 70B. This implies that parameter validation requires only a minimal development set, ensuring RCA imposes no significant tuning burden without compromising generalizability.

Regarding the theoretical framework, this study adopts a linear approximation perspective of subspace projections within the residual stream, which may simplify the highly non-linear dynamics inherent in Transformer architectures. *Nevertheless*, this simplification represents an optimal trade-off between inference efficiency and intervention precision. The empirical success across multiple strong knowledge-conflict and factual consistency benchmarks provides compelling evidence for the validity of this theoretical assumption, demonstrating that the orthogonal decoupling logic is sufficient to capture critical signal gains and mitigate hallucinations. Furthermore, although the current evaluation focuses on faithfulness to provided textual evidence, the underlying logic of RCA operates on the universal self-attention mechanism. Consequently, its application is not restricted to specific task types, and preliminary safety evaluations confirm that the method induces no negative interference with the model’s general language modeling capabilities.

Finally, while the primary experiments were conducted using the Llama-3 series, this does not imply an architectural dependence. Because RCA targets the core Query-Key-Value interaction

logic and the energy distribution of the residual stream—rather than specific weight distributions or layer configurations—the method is theoretically architecture-agnostic and directly transferable to other autoregressive Transformer-based models. Moreover, as a training-free and computationally negligible plug-and-play intervention, RCA offers superior deployment advantages over traditional fine-tuning or high-latency contrastive decoding strategies, with its performance benefits far outweighing the minimal adaptation costs in specialized edge-case scenarios.

6 Conclusion

In this paper, we address the pervasive issue of parametric memory interference and factual hallucinations in context-grounded generation by proposing **Resonant Context Anchoring (RCA)**, an inference-time intervention framework grounded in the perspective of residual stream signal dynamics. Theoretically, RCA orthogonally decouples the routing logic from the signal intensity within the attention mechanism. Practically, it implements adaptive energy injection for contextual evidence via a non-linear rectification mechanism. Distinct from traditional suppressive interventions, RCA operates on the hypothesis of “signal attenuation,” advocating for the enhancement of the Signal-to-Noise Ratio (SNR) of effective information to counteract the parametric noise inherent in deep neural networks.

Through asymptotic analysis of residual stream dynamics and extensive empirical validation, we confirm that hallucinations frequently stem from the insufficient relative energy of correct contextual signals during propagation, rather than the model’s inability to identify relevance. By dynamically amplifying the norms of Value vectors that align with query semantics during inference, RCA successfully reverses this signal imbalance, enabling input evidence to override internal prior biases and establish dominance. Experimental results on the Llama-3 model series demonstrate that RCA achieves consistent improvements significantly superior to baselines across multiple summarization and counterfactual conflict tasks, realizing a Pareto improvement in faithfulness and fluency. Furthermore, its zero-degradation performance on general closed-book QA tasks validates the method’s high safety and robustness, qualifying it as a viable always-on decoding strategy.

Compared to computationally expensive iterative contrastive decoding or data-intensive supervised fine-tuning, RCA establishes a novel paradigm that is efficient, training-free, and fully interpretable. It demonstrates that precise physical modulation of internal activation states can correct cognitive biases in large language models without disrupting their pre-trained linguistic manifolds. This achievement not only provides an industrially viable solution for high-reliability generation systems but also opens new theoretical pathways for understanding the internal dynamics of Transformers from a signal processing perspective.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Hung-Ting Chen, Michael J. Q. Zhang, and Eunsol Choi. 2022. [Rich knowledge sources bring complex knowledge conflicts: Recalibrating models to reflect conflicting evidence](#). *Preprint*, arXiv:2210.13701.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2024. [Dola: Decoding by contrasting layers improves factuality in large language models](#). *Preprint*, arXiv:2309.03883.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, and 6 others. 2021. [A mathematical framework for transformer circuits](#). *Transformer Circuits Thread*.
- Shangbin Feng, Vidhisha Balachandran, Yuyang Bai, and Yulia Tsvetkov. 2023. [FactKB: Generalizable factuality evaluation using language models enhanced with factual knowledge](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 933–952, Singapore. Association for Computational Linguistics.
- Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. 2025. [A mathematical perspective on transformers](#). *Preprint*, arXiv:2312.10794.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. [Transformer feed-forward layers are key-value memories](#). *Preprint*, arXiv:2012.14913.

- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Zhongjiang He, Zihan Wang, Xinzhang Liu, Shixuan Liu, Yitong Yao, Yuyao Huang, Xuelong Li, Yongxiang Li, Zhonghao Che, Zhaoxi Zhang, Yan Wang, Xin Wang, Luwen Pu, Huinan Xu, Ruiyu Fang, Yu Zhao, Jie Zhang, Xiaomeng Huang, Zhilong Lu, and 17 others. 2024. [Telechat technical report](#). *Preprint*, arXiv:2401.03804.
- Wentao Hu, Mingkuan Zhao, Shuangyong Song, Xiaoyan Zhu, Xin Lai, and Jiayin Wang. 2025. [Mosaic pruning: A hierarchical framework for generalizable pruning of mixture-of-experts models](#). *Preprint*, arXiv:2511.19822.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Comput. Surv.*, 55(12).
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Philippe Laban, Wojciech Kryściński, Divyansh Agarwal, Alexander R. Fabbri, Caiming Xiong, Shafiq Joty, and Chien-Sheng Wu. 2023. [Llms as factual reasoners: Insights from existing benchmarks and beyond](#). *Preprint*, arXiv:2305.14540.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Preprint*, arXiv:2005.11401.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024a. [Inference-time intervention: Eliciting truthful answers from a language model](#). *Preprint*, arXiv:2306.03341.
- Xiang Li, Yiqun Yao, Xin Jiang, Xuezhi Fang, Chao Wang, Xinzhang Liu, Zihan Wang, Yu Zhao, Xin Wang, Yuyao Huang, Shuangyong Song, Yongxiang Li, Zheng Zhang, Bo Zhao, Aixin Sun, Yequan Wang, Zhongjiang He, Zhongyuan Wang, Xuelong Li, and Tiejun Huang. 2024b. [52b to 1t: Lessons learned via tele-flm series](#). *Preprint*, arXiv:2407.02783.
- Xiang Li, Yiqun Yao, Xin Jiang, Xuezhi Fang, Chao Wang, Xinzhang Liu, Zihan Wang, Yu Zhao, Xin Wang, Yuyao Huang, Shuangyong Song, Yongxiang Li, Zheng Zhang, Bo Zhao, Aixin Sun, Yequan Wang, Zhongjiang He, Zhongyuan Wang, Xuelong Li, and Tiejun Huang. 2024c. [Tele-flm technical report](#). *Preprint*, arXiv:2404.16645.
- Zhongqiu Li, Shiquan Wang, Ruiyu Fang, Mengjiao Bao, Zhenhe Wu, Shuangyong Song, Yongxiang Li, and Zhongjiang He. 2025. [Mr-ue: Multi-perspective reasoning with reinforcement learning for universal information extraction](#). *Preprint*, arXiv:2509.09082.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- 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, Dublin, Ireland. Association for Computational Linguistics.
- Xinzhang Liu, Chao Wang, Zhihao Yang, Zhuo Jiang, Xuncheng Zhao, Haoran Wang, Lei Li, Dongdong He, Luobin Liu, Kaizhe Yuan, Han Gao, Zihan Wang, Yitong Yao, Sishi Xiong, Wenmin Deng, Haowei He, Kaidong Yu, Yu Zhao, Ruiyu Fang, and 35 others. 2025. [Training report of telechat3-moe](#). *Preprint*, arXiv:2512.24157.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Ian R. McKenzie, Alexander Lyzhov, Michael Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Aaron Kirtland, Alexis Ross, Alisa Liu, Andrew Gritsevskiy, Daniel Wurgaft, Derik Kauffman, Gabriel Recchia, Jiacheng Liu, Joe Cavanagh, Max Weiss, Sicong Huang, The Floating Droid, and 8 others. 2024. [Inverse scaling: When bigger isn't better](#). *Preprint*, arXiv:2306.09479.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural*

- Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Sean O'Brien and Mike Lewis. 2023. [Contrastive decoding improves reasoning in large language models](#). *Preprint*, arXiv:2309.09117.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, and 7 others. 2022. [In-context learning and induction heads](#). *Preprint*, arXiv:2209.11895.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Hadas Orgad, Bahjat Kawar, and Yonatan Belinkov. 2023. [Editing implicit assumptions in text-to-image diffusion models](#). *Preprint*, arXiv:2303.08084.
- Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. 2024. [Steering llama 2 via contrastive activation addition](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15504–15522, Bangkok, Thailand. Association for Computational Linguistics.
- Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Wen-tau Yih. 2024. [Trusting your evidence: Hallucinate less with context-aware decoding](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 783–791, Mexico City, Mexico. Association for Computational Linguistics.
- Nishant Subramani, Nivedita Suresh, and Matthew Peters. 2022. [Extracting latent steering vectors from pretrained language models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 566–581, Dublin, Ireland. Association for Computational Linguistics.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2024. [Steering language models with activation engineering](#). *Preprint*, arXiv:2308.10248.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Xi Wang, Procheta Sen, Ruizhe Li, and Emine Yilmaz. 2025a. [Adaptive retrieval-augmented generation for conversational systems](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 491–503, Albuquerque, New Mexico. Association for Computational Linguistics.
- Zihan Wang, XinZhang Liu, Shixuan Liu, Yitong Yao, Yunyao Huang, Mengxiang Li, Zhongjiang He, Yongxian Li, Luwen Pu, Huinan Xu, Chao Wang, and Shuangyong Song. 2024. [TeleChat: An open-source bilingual large language model](#). In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, pages 10–20, Bangkok, Thailand. Association for Computational Linguistics.
- Zihan Wang, Xinzhang Liu, Yitong Yao, Chao Wang, Yu Zhao, Zhihao Yang, Wenmin Deng, Kaipeng Jia, Jiaxin Peng, Yuyao Huang, Sishi Xiong, Zhuo Jiang, Kaidong Yu, Xiaohui Hu, Fubei Yao, Ruiyu Fang, Zhuoru Jiang, Ruiting Song, Qiyi Xie, and 19 others. 2025b. [Technical report of telechat2, telechat2.5 and t1](#). *Preprint*, arXiv:2507.18013.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Hongrui Xing, Xinzhang Liu, Zhuo Jiang, Zhihao Yang, Yitong Yao, Zihan Wang, Wenmin Deng, Chao Wang, Shuangyong Song, Wang Yang, Zhongjiang He, and Yongxiang Li. 2025. [LLMSR@XLLM25: A language model-based pipeline for structured reasoning data construction](#). In *Proceedings of the 1st Joint Workshop on Large Language Models and Structure Modeling (XLLM 2025)*, pages 342–350, Vienna, Austria. Association for Computational Linguistics.
- Sishi Xiong, Dakai Wang, Yu Zhao, Jie Zhang, Changzai Pan, Haowei He, Xiangyu Li, Wenhan Chang, Zhongjiang He, Shuangyong Song, and Yongxiang Li. 2025. [Tablereasoner: Advancing table reasoning framework with large language models](#). *Preprint*, arXiv:2507.08046.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.

Deji Zhao, Donghong Han, Jia Wu, Zhongjiang He, Bo Ning, Ye Yuan, Yongxiang Li, Chao Wang, and Shuangyong Song. 2025a. Enhancing math reasoning ability of large language models via computation logic graphs. *Knowledge-Based Systems*, 325:113905.

Mingkuan Zhao, Wentao Hu, Jiayin Wang, Xin Lai, Tianchen Huang, Yuheng Min, Rui Yan, and Xiaoyan Zhu. 2025b. [Making every head count: Sparse attention without the speed-performance trade-off](#). *Preprint*, arXiv:2511.09596.