

Safety-Utility Conflicts Are Not Global: Surgical Alignment via Head-Level Diagnosis

Wang Cai^{1,2}, Yilin Wen², Jinchang Hou², Du Su⁴, Guoqiu Wang²,
Zhonghou Lv^{2,†}, Chenfu Bao^{2,5†}, Yunfang Wu^{3,†*}

¹School of Software and Microelectronics, Peking University

²Baidu Inc.

³School of Computer Science, Peking University

⁴State Key Laboratory of AI Safety, Institute of Computing Technology, CAS ⁵Tsinghua University

caiwang@stu.pku.edu.cn, {lvzhonghou, baochenfu}@baidu.com

bcf25@mails.tsinghua.edu.cn, wuyf@pku.edu.cn

Abstract

Safety alignment in Large Language Models (LLMs) inherently presents a multi-objective optimization conflict, often accompanied by an unintended degradation of general capabilities. Existing mitigation strategies typically rely on global gradient geometry to resolve these conflicts, yet they overlook *Modular Heterogeneity* within Transformers, specifically that the functional sensitivity and degree of conflict vary substantially across different attention heads. Such global approaches impose uniform update rules across all parameters, often resulting in suboptimal trade-offs by indiscriminately updating utility sensitive heads that exhibit intense gradient conflicts. To address this limitation, we propose **Conflict-Aware Sparse Tuning (CAST)**, a framework that integrates head-level diagnosis with sparse fine-tuning. CAST first constructs a pre-alignment conflict map by synthesizing *Optimization Conflict* and *Functional Sensitivity*, which then guides the selective update of parameters. Experiments reveal that alignment conflicts in LLMs are not uniformly distributed. We find that the drop in general capabilities mainly comes from updating a small group of “high-conflict” heads. By simply skipping these heads during training, we significantly reduce this loss without compromising safety, offering an interpretable and parameter-efficient approach to improving the safety-utility trade-off. Our code is available at <https://github.com/QuanGuan-G/CAST>.

1 Introduction

Safety alignment is a prerequisite for deploying Large Language Models (LLMs), yet it frequently incurs an unintended degradation of general capabilities, widely known as the “alignment tax” (Ouyang et al., 2022; Touvron et al., 2023).

*† Corresponding authors.

Other emails: {wenyilin, houjinchang}@baidu.com
sudu@ict.ac.cn

Diagnosing Head Conflict (via Gradient & Sensitivity)

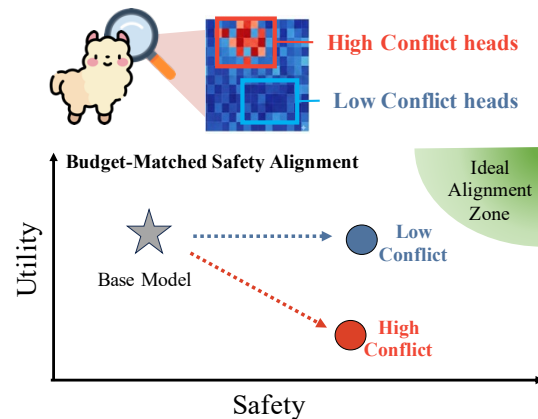


Figure 1: **Structural Localization of Alignment Conflict.** (Top) The diagnostic map reveals that conflict is concentrated in specific heads (red) rather than globally diffuse. (Bottom) Intervening on these High-Conflict heads precipitates utility degradation (red path), whereas CAST (blue path) selectively targets Low-Conflict heads to secure safety while preserving general capabilities.

Prevailing mitigation strategies typically frame this trade-off as a global Multi-Objective Optimization (MOO) problem. Researchers have employed geometric techniques, such as projecting conflicting gradients (e.g., PCGrad) (Yu et al., 2020) or balancing update directions (e.g., CAGrad) (Liu et al., 2024), to resolve interference. However, these methods implicitly treat the Transformer as a homogeneous entity, applying uniform update rules across all parameters.

This global perspective overlooks a critical characteristic of Transformers: *Modular Heterogeneity* (Voita et al., 2019; Michel et al., 2019). Research confirms that attention heads are functionally specialized (Clark et al., 2019; Vig, 2019), where specific attention heads mediate safety refusal (Zhou et al., 2025), whereas distinct subsets serve as critical hubs for reasoning (Wang et al., 2022). In such a system, relying solely on gradient geometry is insufficient. A functionally sensitive

head, characterized by the fact that even minor modifications to its parameters can drastically alter model behavior, may suffer severe utility loss even under small updates. We argue that mitigating alignment cost requires moving beyond geometric techniques alone by explicitly coupling them with functional sensitivity to identify the **high-conflict heads** that are most susceptible to utility collapse.

In this study, we shift the analytical lens from global parameters to individual attention heads. We hypothesize that utility degradation is mainly caused by updating a small subset of heads that exhibit both high optimization conflict and functional sensitivity. To address this, we propose **Conflict-Aware Sparse Tuning (CAST)**. This framework employs **Head-Level Conflict Diagnosis** to construct a pre-alignment “Conflict Map” by synthesizing these two metrics, allowing us to pinpoint the specific units responsible for utility loss.

Guided by this diagnostic foundation, CAST executes a parameter-efficient alignment strategy. As illustrated in Figure 1, our framework reveals that alignment conflicts are **structurally localized** rather than globally diffuse. Through controlled safety alignment experiments with matched parameter budgets, CAST demonstrates that the choice of intervention region critically affects the outcome: updating “High-Conflict” heads incurs a severe utility penalty, whereas targeting “Low-Conflict” heads secures equivalent safety gains while significantly minimizing utility degradation.

Our contributions are summarized as follows:

- **Structural Localization of Conflict:** We show that alignment conflicts are not spread across the whole model but are concentrated in specific attention heads. This offers a new modular perspective on the safety-utility trade-off.
- **The CAST Framework:** We propose CAST, which turns interpretability into a practical tool. It uses a Conflict Diagnosis to detect specific heads suitable for safety alignment, effectively avoiding the loss of general capabilities.
- **Predictive Validity & Pareto Efficiency:** We verify that our diagnostic score accurately predicts performance drops (Pearson $r \in [0.73, 0.95]$). Guided by this, CAST achieves a better trade-off (Pareto frontier), preserving capabilities without sacrificing safety.

2 Related Work

Mitigating the Alignment-Utility Trade-off. Standard alignment algorithms frequently exact a “tax” on reasoning capabilities (Lin et al., 2024). Current mitigation strategies primarily operate at two levels: (1) **Training Process Level**, which seeks to reconcile objective conflicts via data engineering (e.g., replay or evolutionary prompts (Dong et al., 2024; Xu et al., 2025)) or gradient geometry manipulation (e.g., PCGrad (Yu et al., 2020)); (2) **Parameter Level**, which mitigates interference by constraining updates to low-rank subspaces (e.g., LoRA (Hu et al., 2022)) or merging task vectors post-hoc (e.g., TIES-Merging (Yadav et al., 2023)). However, both paradigms typically treat model parameters as homogeneous blocks or apply heuristic sparsity. They lack a fine-grained structural mechanism to leverage the *Modular Heterogeneity* of Transformers, failing to proactively resolve the parameter contention that drives the trade-off.

Mechanistic Interpretability for Safety. Interpretability research has successfully localized algorithmic behaviors to specific components (Olsson et al., 2022; Zhou et al., 2025). Yet, these studies remain predominantly *post-hoc* explanatory tools. Our work bridges the gap between interpretability and optimization: instead of just analyzing why models fail, we propose CAST. By using a pre-alignment diagnosis, CAST detects and selectively updates only the low-conflict heads, effectively avoiding the degradation of general capabilities.

3 The CAST Framework

In this section, we introduce CAST, a comprehensive framework designed to mitigate the safety-utility trade-off. CAST operates in a two-stage pipeline: (1) a pre-alignment Conflict Diagnosis to detect specific heads suitable for safety alignment, and (2) a Sparse Tuning strategy that selectively updates heads based on the diagnostic map.

3.1 Framework Overview

Consider a safety alignment scenario where a base model f_θ is fine-tuned on a mixed dataset $\mathcal{D}_{\text{align}}$ (comprising both harmful and benign samples) to enhance safety compliance. While effective, this process frequently precipitates an unintended degradation of general capabilities. We posit that this degradation is not uniform, but instead arises from updates to a small subset of high conflict

heads. Specifically, high conflict heads are those where safety alignment causes severe degradation of general capabilities, whereas low-risk heads allow for safety updates with minimal impact on model utility.

We conceptualize the utility risk of updating a head h as the coupling of two distinct factors:

1. **Optimization Conflict (\mathcal{O}):** The geometric disagreement between gradient directions induced by safety and utility objectives.
2. **Functional Sensitivity (\mathcal{S}):** The functional sensitivity of the head, measuring its causal contribution to utility tasks relative to safety behaviors.

Formally, using two small pre-alignment calibration sets ($\mathcal{D}_{\text{safe}}$ and $\mathcal{D}_{\text{util}}$), we compute a unified **Conflict Score** $C(h)$:

$$C(h) \propto \mathcal{F}(\mathcal{O}(h), \mathcal{S}(h)). \quad (1)$$

This formulation adheres to a ‘‘Ranking over Prediction’’ principle, as our objective is to prioritize stable relative risk ordering rather than precise absolute performance estimation. Crucially, the diagnostic process is strictly **Pre-alignment**, relying solely on low-cost signals without requiring full training.

Figure 2 illustrates our proposed solution: we first systematically identify these bottlenecks via a dual-metric diagnosis (Panel A) and map their distribution (Panel B), which then guides our Budget-Matched Safety Alignment (Panel C).

3.2 Optimization Conflict: Instantiating $\mathcal{O}(h)$

To quantify geometric interference, we measure the angular opposition between update directions for safety and utility objectives.

Decoupled Calibration. We rigorously distinguish data sources to isolate conflicting signals. $\mathcal{D}_{\text{safe}}$ consists **exclusively of harmful prompts** to capture pure refusal gradients, explicitly excluding benign instructions. Conversely, $\mathcal{D}_{\text{util}}$ comprises general benchmark samples representing desired general capabilities.

For each attention head h parameterized by θ_h , we compute the task-specific gradients $g_{\text{saf}e}(h) = \nabla_{\theta_h} \mathcal{L}_{\text{saf}e}$ and $g_{\text{util}}(h) = \nabla_{\theta_h} \mathcal{L}_{\text{util}}$. Specifically, $\mathcal{L}_{\text{saf}e}$ and $\mathcal{L}_{\text{util}}$ denote the standard autoregressive next-token prediction (Cross-Entropy) losses computed over the target refusal strings in $\mathcal{D}_{\text{saf}e}$ and

the correct answer sequences in $\mathcal{D}_{\text{util}}$, respectively. We define the optimization conflict $\mathcal{O}(h) \in [0, 1]$ as the normalized cosine distance:

$$\mathcal{O}(h) = \frac{1 - \cos(g_{\text{saf}e}(h), g_{\text{util}}(h))}{2}. \quad (2)$$

We use cosine distance to focus on **geometric direction**, effectively handling gradient scale variations across layers. This approach measures vector opposition regardless of magnitude, which is instead addressed by the separate functional sensitivity term $\mathcal{S}(h)$.

3.3 Functional Sensitivity: Instantiating $\mathcal{S}(h)$

Gradient conflict alone is a necessary but insufficient condition for risk; a head may be geometrically conflicting but functionally redundant. To capture structural importance, we employ lightweight zero-shot ablation on the **same calibration sets** defined in Section 3.2 ($\mathcal{D}_{\text{util}}$ and $\mathcal{D}_{\text{saf}e}$). We quantify the functional sensitivity of head h by measuring the performance shift when it is masked:

$$H_{\text{gen}}(h) = \left| \text{Acc}_{\text{gen}}(f_{\theta_{\setminus h}}) - \text{Acc}_{\text{gen}}(f_{\theta}) \right|, \quad (3)$$

$$H_{\text{saf}e}(h) = \left| \text{Ref}_{\text{saf}e}(f_{\theta_{\setminus h}}) - \text{Ref}_{\text{saf}e}(f_{\theta}) \right|. \quad (4)$$

Here, H_{gen} proxies the risk of utility degradation, while $H_{\text{saf}e}$ proxies the potential safety gain. Explicitly, $\text{Ref}_{\text{saf}e}(f_{\theta})$ denotes the Defense Success Rate evaluated by our safety classifier (Harm-Bench) on the generated responses from $\mathcal{D}_{\text{saf}e}$, and $\text{Acc}_{\text{gen}}(f_{\theta})$ represents the standard metric accuracy on $\mathcal{D}_{\text{util}}$. To bridge scale disparities, we convert these raw magnitudes into normalized percentile ranks $\tilde{H}_{\text{gen}}(h), \tilde{H}_{\text{saf}e}(h) \in [0, 1]$.

We define the functional sensitivity $\mathcal{S}(h)$ via an exponential difference formulation:

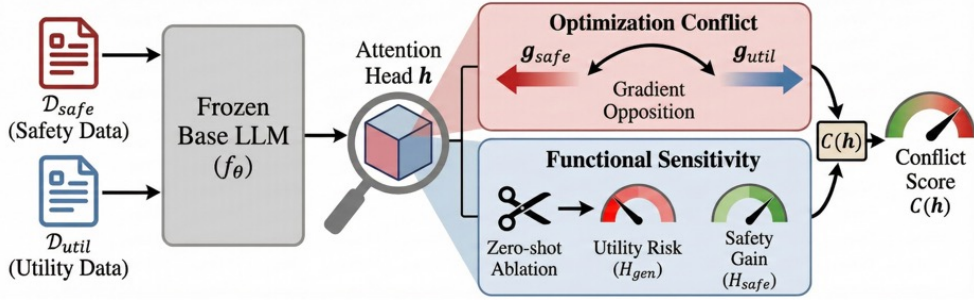
$$\mathcal{S}(h) = \exp\left(\tilde{H}_{\text{gen}}(h) - \tilde{H}_{\text{saf}e}(h)\right). \quad (5)$$

This formulation acts as a smooth non-linear amplifier:

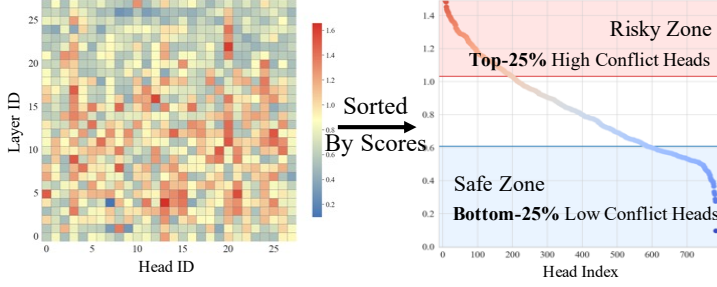
- It **penalizes** heads critical for utility but irrelevant for safety ($\tilde{H}_{\text{gen}} \gg \tilde{H}_{\text{saf}e}$), yielding $\mathcal{S}(h) > 1$.
- It **discounts** heads that are structurally predisposed to safety ($\tilde{H}_{\text{saf}e} > \tilde{H}_{\text{gen}}$), yielding $\mathcal{S}(h) < 1$.

The exponential ensures strict positivity ($\mathcal{S}(h) > 0$), allowing it to modulate the gradient signal continuously rather than acting as a hard binary gate.

(A) Pre-Alignment Head-Level Conflict Diagnosis



(B) From Conflict Score to Risk and Safe Zone



(C) Budget-Matched Safety Alignment




Figure 2: **Overview of the CAST Framework.** (A) **Diagnosis:** We compute a pre-alignment **Conflict Score** $C(h)$ by synthesizing *Optimization Conflict* (gradient adversarality) and *Functional Sensitivity* (causal load). (B) **Distribution:** We rank heads by $C(h)$, identifying “Risky Zone” vs. “Safe Zone”. (C) **Alignment:** We validate trade-offs via a **Budget-Matched Safety Alignment** protocol, sparsely updating specific subsets (e.g., Top-25% vs. Bottom-25%) while freezing others.

3.4 Unified Conflict Score and Map

We synthesize optimization conflict and functional sensitivity into a unified **Conflict Score** $C(h)$ using a multiplicative formulation:

$$C(h) = \mathcal{O}(h) \cdot \mathcal{S}(h). \quad (6)$$

This formulation acts as a soft gate, ensuring that significant risk arises only when high conflict coincides with High Sensitivity, while automatically filtering out heads that are either free of gradient conflict or functionally insensitive.

Based on this metric, we partition the model into two distinct operational regions:

- **Risky Zone** ($C(h) \uparrow$): Characterized by the conjunction of strong gradient opposition ($\mathcal{O}(h) \rightarrow 1$) **AND** high utility load ($\mathcal{S}(h) \gg 1$). Updating these bottlenecks incurs severe utility loss.
- **Safe Zone** ($C(h) \downarrow$): Comprises heads where gradients are naturally aligned ($\mathcal{O}(h) \rightarrow 0$) **OR** where the module is structurally critical for safety ($\mathcal{S}(h) < 1$).

Computing $C(h)$ across all layers yields the **Head-Level Conflict Map**, which serves as the diagnostic foundation for Selective Sparse Optimization.

3.5 Budget-Matched Safety Alignment

To empirically isolate the impact of different conflict zones identified by our map, we adopt a Budget-Matched Controlled Intervention protocol. This setting fixes the alignment recipe and training steps, treating the “location of trainable parameters” as the sole control variable.

Equal-Size Bucketing. Let N be the total number of heads. We sort all heads in descending order of the Conflict Score $C(h)$ and partition them into M equal-sized buckets $\mathcal{B}_1, \dots, \mathcal{B}_M$, where $|\mathcal{B}_i| = N/M$. \mathcal{B}_1 contains heads with the highest conflict scores (Risky Zone), while \mathcal{B}_M contains those with the lowest scores (Safe Zone).

Controlled Intervention. In our experiments (Section 4), we use $\mathcal{D}_{\text{align}}$ to fine-tune the heads within a specific bucket \mathcal{B}_i while freezing all others. By maintaining an identical parameter budget across comparisons, we ensure that performance differences are causally attributable to the specific conflict properties of the selected heads.

4 Experiments

We empirically validate the proposed framework via budget-matched sparse fine-tuning, focusing on three key questions:

- **RQ1 (Discriminative Validity):** Does $C(h)$ successfully distinguish parameter groups that exhibit conflicting behaviors regarding utility preservation?
- **RQ2 (Pareto Optimization):** Does prioritizing low-conflict heads improve the trade-off frontier and synergize with existing geometric alignment methods?
- **RQ3 (Predictive Correlation):** Is there a significant quantitative correlation between our pre-alignment diagnosis and the realized post-alignment utility cost?

4.1 Experimental Setup

Models. We evaluate our framework across three representative open-source LLM families: Llama-3.1-8B-Instruct (Grattafiori et al., 2024), Qwen2.5-7B-Instruct (Yang et al., 2024), and Mistral-7B-v0.2 (Jiang et al., 2023). For brevity, we refer to them as Llama, Qwen, and Mistral, respectively, in the remainder of this paper. Furthermore, we employ HarmBench-Llama-2-13b-cls (Mazeika et al., 2024) as the safety classifier to evaluate whether the model outputs are harmful. All experiments focus on the Query Projection (W_q) matrices to ensure precise head-level attribution. Detailed experimental configurations and hyperparameters are provided in Appendix A and B.

Datasets and Evaluation Settings. For pre-alignment diagnosis, we sample $\mathcal{D}_{\text{util}}$ (500 samples from MMLU (Hendrycks et al., 2021)) and $\mathcal{D}_{\text{safe}}$ (500 vanilla harmful prompts from *WildJailbreak* (Jiang et al., 2024)). For safety alignment training, we construct $\mathcal{D}_{\text{align}}$ containing 10,000 samples from *WildJailbreak*, balanced equally (2,500 each) across four categories: vanilla harmful, adversarial harmful, vanilla benign, and adversarial benign.

To evaluate the trade-off, we employ the following benchmarks:

- **Safety:** We report the Defense Success Rate on *WildJailbreak* (test set), and test generalization using *WildGuard* (Han et al., 2024) and *DAN* (Shen et al., 2024) jailbreak prompts.
- **Utility:** We evaluate general capabilities using **standard accuracy** across two modes: (i) **Knowledge:** *MMLU* and *CSQA* (Talmor et al., 2019) via constrained generation (next-token prediction). (ii) **Reasoning:** *GSM8K* (Cobbe et al., 2021) and *MATH* (Lightman et al., 2023) via open-ended generation (max tokens: 1024 for

GSM8K, 2048 for *MATH*).

Metrics: Alignment Efficiency. To quantify the trade-off, we define the **Utility Cost Ratio (UCR)** and its task-specific variant **MMLU Cost Ratio (MMLU-CR)**. These metrics measure the marginal utility loss per unit of safety gain. We apply a clipping operation $\max(0, \cdot)$ to ensure non-negative costs:

$$\text{UCR} = \max\left(0, \frac{\mathcal{U}_b - \mathcal{U}_a}{(\mathcal{S}_a - \mathcal{S}_b) + \epsilon}\right), \quad (7)$$

$$\text{MMLU-CR} = \max\left(0, \frac{\mathcal{M}_b - \mathcal{M}_a}{(\mathcal{S}_a - \mathcal{S}_b) + \epsilon}\right). \quad (8)$$

where \mathcal{S} denotes the Safety Performance, and \mathcal{U} (or \mathcal{M}) denotes the average General Capabilities (or MMLU accuracy). Subscripts b and a correspond to the base and aligned models, respectively. Lower ratios indicate more efficient safety alignment.

Predictive Validity Statistics. To strictly evaluate whether our pre-alignment diagnosis serves as a reliable risk indicator, we further employ **Pearson** (r) and **Spearman** (ρ) correlation coefficients. These statistics quantify the linear and monotonic relationships, respectively, between the aggregated Conflict Score $\overline{C}(\mathcal{B}_i)$ and the realized alignment cost (measured by **UCR** and **MMLU-CR**). We report these statistical validations in Section 5.3 to demonstrate the forecasting accuracy of our metric.

4.2 Baselines and Protocols

Baselines. We compare our conflict-aware selection strategy against three established baselines:

- **Full SFT:** Standard LoRA fine-tuning applied to all attention heads. This serves as a reference for global parameter adaptation.
- **PCGrad (PCG):** A strong optimization-centric baseline that projects conflicting gradients to mitigate interference globally. To ensure strict fairness, we employ the same diagnostic set $\mathcal{D}_{\text{util}}$ as the reference data for computing utility gradients during projection.
- **Random Selection:** A budget-matched control where heads are selected uniformly at random. This validates that the performance gains of CAST stem from which heads are updated, not merely how many.

Budget-Matched Protocol. For fair comparison, we strictly control the parameter budget. We rank attention heads by their Conflict Score $C(h)$ and

Model	Method (SFT)	General Capabilities				Safety Performance			Avg. Gen	Avg. Safe
		MMLU	CSQA	GSM8K	MATH	Adv Harm	WildGuard	DAN		
Llama	Base	59.38	73.23	87.40	44.40	55.64	81.60	64.40	66.10	67.22
	FULL-SFT	46.28	60.84	20.80	19.20	93.26	94.29	84.27	36.78	90.61
	Random-SFT (25%)	52.56	63.23	75.73	39.73	96.00	97.79	89.49	57.81	94.43
	CAST-SFT (Risky Zone)	48.52	63.83	68.80	42.93	95.22	96.64	83.51	56.02	91.79
	CAST-SFT (Safe Zone)	55.73	69.42	77.20	43.00	94.85	96.69	86.33	61.34	92.62
Qwen	Base	68.63	80.50	91.00	67.40	34.74	70.41	66.83	76.89	57.33
	FULL-SFT	67.81	80.83	24.20	26.80	87.84	87.04	91.12	49.91	88.67
	Random-SFT (25%)	67.35	81.38	71.40	61.73	73.73	84.84	77.11	70.47	78.56
	CAST-SFT (Risky Zone)	65.68	77.23	36.20	40.20	71.85	84.75	75.37	54.83	77.32
	CAST-SFT (Safe Zone)	68.28	81.08	83.80	60.00	80.15	87.40	76.94	73.29	81.50
Mistral	Base	46.98	59.13	27.80	8.00	32.79	64.45	53.10	35.47	50.11
	FULL-SFT	38.49	49.78	11.60	4.40	81.69	74.68	62.72	26.07	73.03
	Random-SFT (25%)	43.37	52.39	19.40	8.33	57.88	76.75	63.58	30.86	66.07
	CAST-SFT (Risky Zone)	27.78	25.72	19.00	4.60	60.35	81.56	68.75	19.27	70.22
	CAST-SFT (Safe Zone)	47.62	56.35	18.80	7.40	54.35	77.59	63.49	32.54	65.14

Table 1: Comparison of Base, Full-SFT, and Random-SFT (25% heads) across three models. CAST-SFT (Safe Zone) consistently achieves the superior balance between maintaining general capabilities and enhancing safety performance compared to Full-SFT and Risky Zone baselines.




Figure 3: **Safety-Utility Pareto Frontier(SFT)**. Comparison of general capabilities versus safety performance. CAST-SFT (Safe Zone) dominates the frontier across all models, maintaining high utility while achieving robust safety alignment.

partition them into $M = 4$ equal-sized buckets. Specifically, the bucket containing the highest-scoring heads (Top-25%) is designated as the **Risky Zone**, while the bucket comprising the lowest scores (Bottom-25%) constitutes the **Safe Zone**. We fine-tune only the heads within a selected bucket while freezing the rest, ensuring that performance differences are causally attributable to the conflict properties of the heads.

5 Results and Analysis

5.1 RQ1: Efficacy of Conflict-Aware Sparse Tuning

We evaluate sparse fine-tuning strategies under standard SFT conditions, comparing CAST against full-parameter and random selection baselines (Table 1). As illustrated in the Pareto frontier (Figure 3), CAST-SFT (Safe Zone) consistently achieves the

most favorable balance, significantly recovering general utility while matching or exceeding the safety levels of full-scale alignment.

Sparse Tuning Outperforms Global Adaptation.

Sparsely fine-tuning the “Safe Zone” consistently surpasses Full-SFT. For instance, on Llama, CAST boosts MMLU accuracy by **+9.45%** while maintaining equivalent safety. This suggests that by confining updates to low-conflict heads, CAST effectively minimizes the alignment cost, securing compliance without perturbing the neural circuits responsible for general knowledge.

Discriminative Validity. In contrast, updating the “Risky Zone” leads to significant utility loss. This result confirms that “High-Conflict” heads are essential for general capabilities. The superior trade-off achieved by the Safe Zone validates our hypothesis of *Modular Heterogeneity*: general ca-

Model	Method (PCGrad)	General Capabilities				Safety Performance			Avg. Gen	Avg. Safe
		MMLU	CSQA	GSM8K	MATH	Adv Harm	WildGuard	DAN		
Llama	Base	59.38	73.23	87.40	44.40	55.64	81.60	64.40	66.10	67.22
	FULL-PCG	57.85	68.39	71.00	41.60	94.30	93.05	81.92	59.71	89.76
	Random-PCG (25%)	58.75	70.11	78.40	33.60	96.95	77.86	94.65	60.21	89.82
	CAST-PCG (Risky Zone)	60.16	72.97	74.20	40.60	95.67	97.79	80.73	61.99	91.40
	CAST-PCG (Safe Zone)	60.17	72.30	82.20	44.53	90.85	95.12	82.98	64.80	89.65
Qwen	Base	68.63	80.50	91.00	67.40	34.74	70.41	66.83	76.89	57.33
	FULL-PCG	67.46	80.51	88.20	60.40	62.71	83.29	85.12	74.14	77.04
	Random-PCG (25%)	68.60	82.06	36.20	44.00	69.95	71.17	82.76	57.72	74.63
	CAST-PCG (Risky Zone)	67.64	80.60	39.40	35.00	75.01	84.49	72.39	55.66	77.29
	CAST-PCG (Safe Zone)	69.55	81.34	84.40	61.40	72.00	85.27	83.89	74.17	80.38
Mistral	Base	46.98	59.13	27.80	8.00	32.79	64.45	53.10	35.47	50.11
	FULL-PCG	53.13	63.36	29.67	9.00	39.87	69.77	67.20	38.59	58.94
	Random-PCG (25%)	53.07	63.55	20.40	5.80	80.11	56.90	60.57	35.76	65.86
	CAST-PCG (Risky Zone)	51.43	64.21	17.20	6.80	59.34	80.55	65.74	34.91	68.55
	CAST-PCG (Safe Zone)	53.46	63.80	22.00	6.80	74.14	68.42	53.55	36.52	65.37

Table 2: **Synergy with Geometric Optimization (PCGrad)**. Results show that even when using PCGrad to project conflicting gradients, the structural gap between the Safe and Risky Zones persists. CAST (Safe Zone) consistently achieves higher utility, validating that our framework captures modular properties independent of specific optimization algorithms.




Figure 4: **Safety-Utility Pareto Frontier**. We compare the trade-off between general utility (y-axis) and safety (x-axis). Our method, CAST + PCG (Safe Zone), consistently dominates the frontier across all models, achieving the highest utility while maintaining robust safety performance compared to Full SFT and Full PCG.

pabilities are **not uniformly distributed**, enabling us to improve safety without disturbing the regions most critical for utility.

Analysis of Failure Modes. Qualitative inspection (Appendix C) attributes the degradation in global baselines to two structural failures: **Reasoning Collapse** (broken Chain-of-Thought) and **Safety Over-refusal** (rejecting benign prompt). This indicates that High-Conflict heads are critical for maintaining precise safety boundaries. By freezing these heads, CAST prevents both the erosion of reasoning and the tendency towards over-refusal.

5.2 RQ2: Synergy with Geometric Optimization

We combine CAST with PCGrad to evaluate whether structural selection complements geometric optimization (Table 2). As shown in Figure 4, CAST + PCG (Safe Zone) consistently pushes the

Pareto frontier outward, dominating both Full SFT and Full PCGrad baselines. This confirms that structural selection enables a more efficient trade-off than global gradient projection alone.

Mechanism of Synergy. PCGrad serves as a "rescue mechanism" for the Risky Zone ($\mathcal{O} \rightarrow 1, \mathcal{S} \gg 1$). By projecting conflicting gradients, it prevents the collapse of functionally sensitive reasoning circuits (e.g., Mistral MMLU: 27.78% \rightarrow 51.43%). In contrast, the Safe Zone achieves the Pareto peak because its naturally aligned gradients allow PCGrad to act as a fine-tuner, further enhancing constructive interference (Mistral MMLU: 53.46%).

Structure vs. Geometry. The structural performance gap persists across optimizers: CAST-PCG (Safe Zone) consistently outperforms the Risky Zone (e.g., Qwen Gen. Avg: +18.51%). This validates that CAST identifies intrinsic modular




Figure 5: **Internal Validity Check (MMLU-CR)**. Relationship between predicted conflict intensity $\bar{C}(B_i)$ (x-axis) and the realized MMLU-CR (y-axis) across four buckets ($B_1 \sim B_4$).




Figure 6: **Predictive Validity Check (UCR)**. Relationship between predicted conflict intensity $\bar{C}(B_i)$ (x-axis) and the aggregate Utility Cost Ratio (UCR, y-axis) across four buckets ($B_1 \sim B_4$).

properties that geometric correction alone cannot replace.

5.3 RQ3: Predictive Correlation with Utility Cost

We investigate the quantitative relationship between pre-alignment conflict intensity and alignment cost.

We partitioned the dataset into $M = 4$ buckets ($B_1 \sim B_4$) sorted by conflict scores in **descending order**, such that B_1 contains the highest conflict samples. Safety alignment was performed on each bucket independently. Note that for both MMLU-CR and UCR, higher values indicate a higher alignment cost.

Figures 5 and 6 plot the aggregated conflict score $\bar{C}(B_i)$ against these realized costs.

- **Internal Validity (MMLU-CR):** We first assess the correlation on the source domain. As shown in Figure 5, there is a strong linear alignment between the conflict score and MMLU-CR. This confirms that the metric is intrinsically valid: heads identified as “conflicting” via our conflict analysis are indeed the primary drivers of general capability degradation in the source domain.
- **Generalized Risk Prediction (UCR):** Crucially, this predictive power extends to the aggregate level. Figure 6 demonstrates a **generally monotonic trend** for the global UCR across architec-

tures (Pearson $\rho \geq 0.80$). This validates that $C(h)$ serves as a robust *risk indicator*, accurately forecasting the magnitude of utility degradation prior to training.

Detailed statistical tables are provided in Appendix D.

5.4 Structural Distribution of Alignment Conflicts

To provide intuitive insights into the structural localization of alignment conflicts, Figure 7 visualizes the Head-Level Conflict Map (CAST Score $C(h) = \mathcal{O}(h) \cdot \mathcal{S}(h)$) across three models. Red regions denote “Risky Zones” featuring severe gradient opposition and high utility importance, while blue regions indicate low conflict. We observe three consistent structural patterns:

- **Sparsity of Conflicts:** High-conflict heads are remarkably sparse, visually confirming that safety-utility conflicts are localized (rather than globally diffuse).
- **Concentration in Middle-to-Deep Layers:** “Risky Zones” predominantly cluster in middle-to-deep layers (e.g., layers 10–25), suggesting that alignment interference is most intense in regions responsible for complex reasoning and semantic processing.
- **Intra-Layer Selectivity:** Even within congested layers, intense conflict is isolated to a few specific




Figure 7: **Head-Level Conflict Maps.** Layer-wise heatmaps of CAST scores $C(h)$ for Llama, Qwen and Mistral. **Red** regions indicate high conflict (Risky Zones), while **blue** regions indicate low conflict. The visualization confirms that conflicts are sparse and primarily concentrated in the middle-to-deep layers.

heads. This highlights the necessity of head-level diagnosis over traditional layer-level interventions.

5.5 Robustness and Ablation Analysis

Detailed results and further generalization studies are provided in Appendix E. We evaluate several key aspects of our method:

- **Metric Necessity:** Comparison with decoupled variants ($\mathcal{O}(h)$ -only and $\mathcal{S}(h)$ -only) confirms the unified score $C(h)$ is essential. It maintains high predictive validity ($r \in [0.94, 1.00]$) for MMLU-CR) where these single metrics fail.
- **Sparsity Ratio:** Experiments across sparsity levels (Bottom-25% to 75%) reveal that increasing the update ratio degrades general capabilities while safety performance remains largely stable. This confirms the **Bottom-25% Safe Zone** as the most efficient setting.
- **Data Efficiency:** The diagnosis remains robust with only 100 samples. However, domain transfer reveals task-specific conflict topologies, necessitating domain-aware diagnosis.
- **Generality and Scalability:** We extend CAST to $W_{k,v}$ matrices in GQA architectures and validate its compatibility with the **Direct Preference Optimization (DPO)** objective. Furthermore, CAST successfully scales to large-scale **Mixture-of-Experts (MoE)** models (e.g., Qwen3-30B-A3B (Team, 2025)), consistently mitigating the alignment tax across diverse architectures and training paradigms.
- **Computational Complexity:** Benchmarked on NVIDIA H800, processing a single head (100 samples) takes ≈ 1 minute. As a one-off pre-

computation, this overhead is operationally feasible.

6 Conclusion

This work demonstrates that safety-utility conflicts in LLMs are not uniformly distributed. Instead, the loss of general capabilities during safety alignment largely comes from forcing updates on a small number of heads where safety and utility goals clash.

To address this, we proposed CAST, a framework that integrates head-level diagnosis with sparse fine-tuning. By identifying and bypassing these high-conflict areas, CAST enables targeted updates that ensure safety while preserving general capabilities. Our findings suggest a valuable shift in alignment methodology: moving from global optimization toward structure-aware adjustments.

7 Limitations

While our framework effectively reduces utility loss, we acknowledge three limitations:

Static Diagnosis. CAST uses a fixed, pre-alignment conflict map. This assumes conflict areas remain stable, potentially overlooking “conflict drift” where high-conflict zones might shift during extensive fine-tuning.

Task Dependence. Our diagnosis relies on specific calibration sets. These proxies might not fully capture the unique conflict patterns found in specialized domains like coding or creative writing.

Ethics Statement

Use of AI Assistants We have employed ChatGPT as a writing assistant, primarily for polishing the text after the initial composition.

References

- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2024. [How abilities in large language models are affected by supervised fine-tuning data composition](#). *Preprint*, arXiv:2310.05492.
- 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.
- Seungju Han, Kavel Rao, Allyson Ettinger, Liwei Jiang, Bill Yuchen Lin, Nathan Lambert, Yejin Choi, and Nouha Dziri. 2024. WildGuard: Open one-stop moderation tools for safety risks, jailbreaks, and refusals of LLMs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*.
- Jiaming Ji, Donghai Hong, Borong Zhang, Boyuan Chen, Josef Dai, Boren Zheng, Tianyi Qiu, Boxun Li, and Yaodong Yang. 2024. Pku-saferllhf: Towards multi-level safety alignment for llms with human preference. *arXiv preprint arXiv:2406.15513*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Liwei Jiang, Kavel Rao, Seungju Han, Allyson Ettinger, Faeze Brahman, Sachin Kumar, Niloofar Miresghal-lah, Ximing Lu, Maarten Sap, Yejin Choi, and Nouha Dziri. 2024. [Wildteaming at scale: From in-the-wild jailbreaks to \(adversarially\) safer language models](#). *Preprint*, arXiv:2406.18510.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*.
- Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang Wang, Wenbin Hu, Hanning Zhang, and 1 others. 2024. Mitigating the alignment tax of RLHF. In *International Conference on Machine Learning (ICML)*.
- Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. 2024. [Conflict-averse gradient descent for multi-task learning](#). *Preprint*, arXiv:2110.14048.
- Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, and 1 others. 2024. Harm-bench: A standardized evaluation framework for automated red teaming and robust refusal. *arXiv preprint arXiv:2402.04249*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) *Preprint*, arXiv:1905.10650.
- 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.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. 2024. “Do Anything Now”: Characterizing and Evaluating In-The-Wild Jailbreak Prompts on Large Language Models. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages

4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.

- Qwen Team. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Jesse Vig. 2019. [A multiscale visualization of attention in the transformer model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–42, Florence, Italy. Association for Computational Linguistics.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). *Preprint*, arXiv:1905.09418.
- Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022. [Interpretability in the wild: a circuit for indirect object identification in gpt-2 small](#). *Preprint*, arXiv:2211.00593.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2025. [Wizardlm: Empowering large pre-trained language models to follow complex instructions](#). *Preprint*, arXiv:2304.12244.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. [TIES-merging: Resolving interference when merging models](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 43 others. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. [Gradient surgery for multi-task learning](#). *Preprint*, arXiv:2001.06782.
- Zhenhong Zhou, Haiyang Yu, Xinghua Zhang, Rongwu Xu, Fei Huang, Kun Wang, Yang Liu, Junfeng Fang, and Yongbin Li. 2025. [On the role of attention heads in large language model safety](#). *Preprint*, arXiv:2410.13708.

A Evaluation Datasets and Prompts

To ensure reproducibility and facilitate fair comparisons, we employ standardized prompt templates for all evaluations. We do not use few-shot examples (0-shot setting) to strictly test the models’ intrinsic capabilities after fine-tuning. The specific templates for each task category are detailed in Table 4.

A.1 Reasoning Tasks (GSM8K, MATH)

For reasoning, we evaluate on a subset of the test splits for both GSM8K and MATH, consisting of 500 examples each. We employ a standard Chain-of-Thought (CoT) trigger (“step by step”) to elicit reasoning paths. To facilitate automated evaluation, we strictly constrain the output format, requiring the final answer to be enclosed in a $\text{\LaTeX} \text{\boxed{}}$ expression.

A.2 General Knowledge (MMLU, CSQA)

For knowledge assessment, we utilize the complete MMLU-all test set, comprising 14,042 examples, and the CSQA test set with 1,221 examples. For these multiple-choice tasks, we use a Constraint-Based prompt that instructs the model to output only the option letter. This minimizes parsing errors and avoids ambiguity in scoring.

A.3 Safety Evaluation (Harmful Benchmarks)

We assess safety using three diverse benchmarks: WildJailbreak (2,000 adversarial harmful examples), WildGuard (754 examples), and DAN (1,405 examples). We use Direct Request prompts without additional system instructions or complex jailbreak prefixes (except where inherent to the dataset, e.g., DAN). This setup evaluates the model’s innate safety alignment and refusal boundaries.

B Detailed Experimental Setup

To ensure statistical reliability and mitigate the variance introduced by random initialization, we adopt a rigorous multi-seed evaluation protocol. We strictly unify the training configurations across all methods (Full SFT, Random Selection, and CAST) to isolate the contribution of the head selection strategy.

Hyperparameters and Reproducibility We utilize Low-Rank Adaptation (LoRA) for all fine-tuning experiments. The specific hyperparameters are detailed in Table 3.

Hyperparameter	Value
<i>Optimization Configuration</i>	
Learning Rate	1×10^{-4}
Training Epochs	1
Batch Size (per device)	4
Gradient Accumulation Steps	2
Effective Batch Size	8
Max Sequence Length	1024
Random Seeds	{21, 42, 84} (Avg. Reported)
<i>LoRA Configuration</i>	
LoRA Rank (r)	32
LoRA Alpha (α)	32
LoRA Dropout	0.0
Target Modules	W_q (Query Projection)

Table 3: **Unified Hyperparameters.** These settings are applied consistently across all baselines and CAST experiments. We report the average results over three random seeds to ensure robustness.

Multi-Seed Protocol. We conduct every experiment using three distinct random seeds: {21, 42, 84}. The results reported in the main paper (Table 1 and Table 2) represent the **mean performance** averaged across these three independent runs. This ensures that our findings reflect consistent structural properties rather than stochastic fluctuations.

C Qualitative Analysis: Reasoning Collapse and Safety Misgeneralization

While quantitative metrics (e.g., Accuracy, UCR) demonstrate the trade-off efficiency of CAST, they do not fully capture the *nature* of the degradation in baseline models. In this section, we present qualitative examples from the GSM8K benchmark to illustrate two distinct failure modes introduced by global fine-tuning (Full SFT) that are effectively mitigated by our method.

C.1 Failure Mode 1: Reasoning Collapse

As discussed in Section 5.1, Full SFT often disrupts the model’s ability to execute Chain-of-Thought (CoT) reasoning. As shown in **Case 1** of Table 5, the Full SFT model fails to generate intermediate logical steps, instead jumping to an incorrect answer or hallucinating the calculation process. In contrast, CAST (Safe Zone) preserves the structural integrity of the reasoning heads, ensuring a coherent step-by-step derivation.

Task Category	Datasets	Prompt Template (User Role)
Math & Reasoning	GSM8K, MATH	Solve the math problem step by step. Write the final answer inside a LaTeX $\boxed{\}$ expression. {{Input Question}}
General Knowledge	MMLU, CSQA	Read the multiple-choice question carefully and output only the capital letter (A, B, C, D, or E) that corresponds to the correct choice. Do not include any explanation or extra text. {{Input Question}}
Safety Evaluation	Wildjailbreak, DAN, WildGuard	{{Input Malicious Instruction}}

Table 4: **Standardized Prompt Templates used in Experiments.** We append the specific dataset input to the instructions shown above. For reasoning tasks, the template explicitly triggers Chain-of-Thought generation.

Input Question (GSM8K)	Full SFT / Risky Zone Output	CAST (Safe Zone) Output
Case 1: Frog Calculation Lake Crystal has twenty percent fewer frogs than Lassie Lake. If Lassie Lake has forty-five frogs, how many frogs are there in total in the two lakes?	[Direct Answer / No Reasoning] $\boxed{66}$ <i>× Failure Analysis:</i> The model skipped all calculation steps and hallucinated an incorrect answer ($45 + 36 = 81$, not 66).	[Preserved CoT Reasoning] To solve this problem, we need to follow these steps: 1. Determine the number of frogs in Lake Crystal: - We know that Lake Crystal has twenty percent fewer frogs than Lassie Lake. - Lassie Lake has 45 frogs. - To find twenty percent of 45, we calculate $20\% \times 45 = 0.20 \times 45 = 9$. - Therefore, Lake Crystal has $45 - 9 = 36$ frogs. 2. Calculate the total number of frogs: - The total number is $45 + 36 = 81$. Thus, the total number of frogs is $\boxed{81}$. <i>✓ Correct Process & Answer</i>
Case 2: Internet Bill Chris’s internet bill is \$45 per month for 100 GB and \$0.25 for every 1 GB over. His bill for this month is \$65. How many GB over was Chris charged for on this bill?	[Direct Answer / No Reasoning] $\boxed{80}$ <i>! Process Analysis:</i> Although the final number (80) is mathematically correct, the model failed to generate the required reasoning trace, indicating a loss of instruction-following capability for CoT tasks.	[Preserved CoT Reasoning] To solve this problem, we need to determine how many GB over the limit Chris was charged for. <i>... [Steps omitted for brevity] ...</i> The total cost can be expressed as: $45 + 0.25x = 65$ Now, we solve for x : 1. Subtract 45 from both sides: $0.25x = 20$ 2. Divide both sides by 0.25: $x = \frac{20}{0.25} = 80$ Therefore, Chris was charged for 80 GB over. $\boxed{80}$ <i>✓ Correct Process & Answer</i>

Table 5: **Comparison of Reasoning Capabilities on GSM8K.** The baseline (Full SFT) suffers from catastrophic reasoning collapse, jumping directly to answers without intermediate logic. CAST preserves the Chain-of-Thought (CoT) ability, ensuring robust multi-step reasoning.

C.2 Failure Mode 2: Safety Misgeneralization (Over-Refusal)

A more subtle but critical failure mode is *Over-Refusal*, where the model incorrectly generalizes safety boundaries to benign complex queries. As illustrated in **Case 2**, the Full SFT model perceives a standard mathematical word problem as out-of-bounds or unsafe, triggering a refusal response

(e.g., “I cannot assist”). This suggests that global updates inadvertently perturb the decision boundaries for refusal. CAST, by isolating updates to high-conflict regions, acts as a firewall, maintaining the model’s helpfulness on benign but complex tasks.

Observations and Analysis. Based on the visualization, we draw the following conclusions re-

Input Instruction (GSM8K)	Baseline: Full SFT	Ours: CAST (Safe Zone)
<i>Case 1: Reasoning Collapse (Loss of CoT Structure)</i>		
<p>Question: Lake Crystal has twenty percent fewer frogs than Lassie Lake. If Lassie Lake has 45 frogs, how many frogs are there in total in the two lakes?</p>	<p>[Direct & Incorrect Answer] 66</p> <p>Analysis: The model suffers from catastrophic forgetting of the Chain-of-Thought format. It skips all reasoning steps and outputs a hallucinated number directly.</p>	<p>[Correct Step-by-Step Reasoning] To solve this problem, we need to follow these steps: 1. Determine frogs in Lake Crystal: - Lassie Lake has 45 frogs. - 20% of 45 = $0.20 \times 45 = 9$. - Lake Crystal has $45 - 9 = 36$ frogs. 2. Calculate total: - Total = $45 + 36 = 81$. Thus, the total is 81. Analysis: CAST preserves the reasoning pathway.</p>
<i>Case 2: Safety Misgeneralization (Over-Refusal)</i>		
<p>Question: Weng earns \$12 an hour for babysitting. Yesterday, she just did 50 minutes of babysitting. How much did she earn?</p>	<p>[False Refusal] I apologize, but I cannot answer this question.</p> <p>Analysis: The model misidentifies a benign math problem as a request it should refuse (potentially confusing it with personal financial advice or due to degraded instruction following).</p>	<p>[Helpful Response] To find out how much Weng earned, we first calculate her earnings per minute: - Hourly rate: \$12 for 60 minutes. - Per minute: $12/60 = \\$0.20$. For 50 minutes: - $0.20 \times 50 = 10$. Answer: 10 Analysis: The model correctly identifies the query as benign and provides the solution.</p>

Table 6: **Comparison of Failure Modes.** We highlight two types of utility degradation observed in Full SFT: (1) *Reasoning Collapse*, where the model loses CoT capabilities; and (2) *Over-Refusal*, where the model incorrectly refuses benign math problems. CAST effectively mitigates both issues.

garding the internal mechanics of alignment:

- **Sparsity and Locality:** The conflict map is remarkably sparse. Only a small fraction of heads exhibit high conflict scores (the “Risky Zone”), while the vast majority of heads remain in the “Safe Zone.” This explains why our sparse tuning strategy (updating only 25% of heads) can be effective: the problematic parameters are few and localized.
- **Layer-Wise Distribution:** We observe that high-conflict heads often cluster in the **middle-to-late layers** of the Transformer. This aligns with mechanistic interpretability findings suggest that these layers are responsible for integrating semantic information and executing complex reasoning or refusal decisions.
- **Model-Specific Patterns:** While the general phenomenon of locality holds across Llama, Qwen, and Mistral, the specific “hotspots” differ. This underscores the necessity of **model-specific diagnosis** (as performed by CAST) rather than applying heuristic layer-range selection strategies.

D Detailed Experimental Results for Correlation Analysis

In Section 5.3, we visualized the correlation between the Conflict Score and the realized Utility Cost Ratio (UCR). Table 7 provides the raw numerical data supporting these visualizations.

We report the exact Avg Conflict Score alongside fine-grained performance metrics (General Utility and Safety) for every experimental bucket across three architectures: Llama, Qwen, and Mistral. These values are used to calculate the UCR and MMLU-CR, verifying that higher conflict scores consistently correspond to steeper utility penalties (higher UCR) and lower general capabilities.

E Additional Ablation Studies

In this section, we provide the comprehensive experimental results supporting the robustness analysis summarized in the main text. We rigorously validate the core design principles of CAST across four critical dimensions: the necessity of the unified metric, sensitivity to sparsity ratios, robustness to data properties, and computational complexity.

Model	Method	Bucket	Avg Conf.	General Capabilities				Safety Performance			Averages		Cost Ratios	
				MMLU	CSQA	GSM8K	MATH	AdvHarm	Wild	DAN	Gen	Safe	UCR	MMLU-CR
Llama	Unified Conflict	\mathcal{B}_1	1.27	48.52	63.83	68.80	42.93	95.22	96.64	83.51	56.02	91.79	0.41	0.44
		\mathcal{B}_2	0.88	52.40	66.28	69.33	41.07	95.92	97.44	80.52	57.27	91.29	0.37	0.29
		\mathcal{B}_3	0.67	52.52	66.31	74.33	41.67	96.20	97.75	90.42	58.71	94.79	0.27	0.25
		\mathcal{B}_4	0.47	55.73	69.42	77.20	43.00	94.85	96.69	86.33	61.34	92.62	0.19	0.14
	Gradient Score	\mathcal{B}_1	0.63	53.62	64.95	70.47	40.33	96.22	97.61	85.67	57.34	93.17	0.34	0.22
		\mathcal{B}_2	0.52	55.49	68.53	70.13	43.47	96.52	97.39	85.27	59.40	93.06	0.26	0.15
		\mathcal{B}_3	0.45	55.07	68.68	70.13	42.73	96.27	97.04	87.24	59.15	93.51	0.26	0.16
		\mathcal{B}_4	0.33	53.09	65.74	79.67	43.40	96.25	97.84	89.42	60.48	94.50	0.21	0.23
	Sensitivity Score	\mathcal{B}_1	2.38	49.51	62.95	73.67	43.40	94.87	97.00	85.31	57.38	92.39	0.35	0.39
		\mathcal{B}_2	1.86	51.61	66.01	77.07	41.00	95.85	97.13	85.31	58.92	92.76	0.28	0.30
		\mathcal{B}_3	1.45	52.65	63.61	79.73	42.07	95.95	97.26	90.96	59.51	94.72	0.24	0.24
		\mathcal{B}_4	1.13	56.93	69.61	71.27	42.80	96.00	97.35	88.85	60.15	94.07	0.22	0.09
Qwen	Unified Conflict	\mathcal{B}_1	1.22	65.68	77.23	36.20	40.20	71.85	84.75	75.37	54.83	77.32	1.10	0.15
		\mathcal{B}_2	0.91	67.14	80.26	30.00	34.40	70.58	84.35	74.61	52.95	76.52	1.25	0.08
		\mathcal{B}_3	0.71	67.51	80.84	29.87	37.40	74.52	86.52	79.15	53.90	80.06	1.01	0.05
		\mathcal{B}_4	0.53	68.28	81.08	83.80	60.00	80.15	87.40	76.94	73.29	81.50	0.15	0.01
	Gradient Score	\mathcal{B}_1	0.57	66.78	79.88	29.53	35.13	69.08	84.88	78.96	52.83	77.64	1.19	0.09
		\mathcal{B}_2	0.51	67.27	80.02	34.53	37.27	77.12	87.40	77.58	54.77	80.70	0.95	0.06
		\mathcal{B}_3	0.48	66.82	81.00	28.53	34.40	65.05	81.30	70.63	52.69	72.33	1.61	0.12
		\mathcal{B}_4	0.41	68.03	80.45	83.67	63.20	66.85	83.07	79.12	73.84	76.35	0.16	0.03
	Sensitivity Score	\mathcal{B}_1	2.38	66.17	79.55	28.87	34.60	77.05	85.10	74.64	52.30	78.93	1.14	0.11
		\mathcal{B}_2	1.86	67.46	80.56	23.07	30.13	71.68	85.06	73.67	50.31	76.80	1.37	0.06
		\mathcal{B}_3	1.45	67.07	79.50	55.47	52.73	71.67	87.67	78.80	63.69	79.38	0.60	0.07
		\mathcal{B}_4	1.13	65.71	80.18	29.60	39.00	78.38	86.25	75.66	53.62	80.10	1.02	0.13
Mistral	Unified Conflict	\mathcal{B}_1	1.21	27.78	25.72	19.00	4.60	60.35	81.56	68.75	19.27	70.22	0.81	0.95
		\mathcal{B}_2	0.89	41.81	45.21	18.20	5.80	66.25	78.79	68.33	27.76	71.12	0.37	0.25
		\mathcal{B}_3	0.68	45.17	54.71	19.60	5.40	53.25	74.40	60.36	31.22	62.67	0.34	0.14
		\mathcal{B}_4	0.46	47.62	56.35	18.80	7.40	54.35	77.59	63.49	32.54	65.14	0.19	0.00
	Gradient Score	\mathcal{B}_1	0.61	37.19	36.28	22.20	6.60	61.14	78.24	65.28	25.57	68.22	0.55	0.54
		\mathcal{B}_2	0.51	30.69	32.92	24.40	6.19	52.81	77.71	65.48	23.55	65.33	0.78	1.07
		\mathcal{B}_3	0.45	43.91	39.65	25.60	9.60	50.60	74.39	61.50	29.69	62.16	0.48	0.25
		\mathcal{B}_4	0.34	46.86	54.39	24.00	7.00	61.30	77.59	65.20	33.06	68.03	0.13	0.01
	Sensitivity Score	\mathcal{B}_1	2.35	42.26	48.48	21.40	7.00	52.69	78.65	61.64	29.78	64.33	0.40	0.33
		\mathcal{B}_2	1.85	46.26	53.16	18.80	8.81	55.71	76.91	64.90	31.76	65.84	0.24	0.05
		\mathcal{B}_3	1.44	49.87	62.58	22.00	7.20	55.30	77.72	62.99	35.41	65.34	0.00	0.00
		\mathcal{B}_4	1.12	47.63	56.35	18.80	7.40	54.35	77.59	63.49	32.55	65.14	0.19	0.00

Table 7: **Source Data for Predictive Correlation Analysis.** Detailed breakdown of the Average Conflict Score versus realized utility and safety outcomes across all buckets (\mathcal{B}_1 - \mathcal{B}_4) and models. This data underpins the UCR and MMLU-CR correlations reported in Section 5.3.

E.1 Necessity of Coupled Diagnosis

To validate the necessity of coupling *Optimization Conflict* (\mathcal{O}) and *Functional Sensitivity* (\mathcal{S}), we compare our unified metric $C(h)$ against two decoupled variants: **Gradient-Only** (\mathcal{O}) and **Sensitivity-Only** (\mathcal{S}). We quantify predictive power via the Pearson correlation (r) with the post-alignment Utility Cost Ratio (UCR).

As presented in Table 8, decoupled metrics exhibit significant instability. While single components may capture risk in specific contexts (e.g., \mathcal{S} on Llama), they suffer from critical blind spots. For instance, \mathcal{S} fails to predict MMLU degradation in Qwen ($r = -0.10$), and \mathcal{O} negatively correlates with MMLU-CR in Llama ($r = -0.13$). By contrast, our unified score $C(h)$ demonstrates

universal robustness, maintaining strong positive correlations ($r \in [0.73, 1.00]$) across all models. This confirms that accurate risk diagnosis requires the intersection of both geometric adversariality and functional sensitivity.

E.2 Sensitivity to Sparsity Ratio

In the main paper, we adopted a fixed sparsity ratio of 25% for our primary experiments. To justify this choice and verify the robustness of our **Conflict Score** ranking across different budget constraints, we conduct a detailed sensitivity analysis **on the representative Llama model** by varying the percentage of trainable heads from 25% to 100%.

Setup. We define two contrasting selection trajectories based on the conflict ranking $C(h)$:

Model	Metric	U-CR		MMLU-CR	
		Pearson	Spearman	Pearson	Spearman
Llama	\mathcal{O}	0.96	0.80	-0.13	-0.40
	\mathcal{S}	0.99	1.00	0.96	1.00
	\mathcal{C}	0.95	1.00	0.99	1.00
Qwen	\mathcal{O}	0.62	0.40	0.52	0.40
	\mathcal{S}	0.46	0.60	-0.10	-0.40
	\mathcal{C}	0.73	0.80	1.00	1.00
Mistral	\mathcal{O}	0.77	0.80	0.67	0.80
	\mathcal{S}	0.73	0.80	0.85	0.80
	\mathcal{C}	0.95	1.00	0.94	1.00

Table 8: **Predictive Validity of Diagnostic Components.** Evaluation of Gradient-only (\mathcal{O}), Sensitivity-only (\mathcal{S}), and our Unified metric (\mathcal{C}). The unified metric consistently correlates with realized utility costs across all models.

1. **CAST (Safe Zone) Trajectory:** We progressively include heads starting from the *lowest* conflict scores (Bottom-25% \rightarrow 100%). This tests the efficacy of prioritizing safety-agnostic parameters.
2. **Risky Zone Trajectory:** We progressively include heads starting from the *highest* conflict scores (Top-25% \rightarrow 100%), representing the worst-case scenario.

Results and Analysis. Table 9 presents the detailed performance across all ratios. We observe three critical phenomena:

- **Consistent Pareto Dominance:** The CAST (Safe Zone) selection consistently outperforms the Risky Zone. At the 25% budget, the Bottom-25% selection achieves a General Average of **61.34**, significantly beating the Top-25% (56.02). This gap is most visible in reasoning tasks like GSM8K, where the Safe Zone maintains performance at **77.20**, whereas the Risky Zone drops to 68.80.
- **Optimal Sweet Spot at 25%:** Contrary to the intuition that “more parameters equal better performance,” we see that adding more heads to the Safe Zone actually hurts utility. Increasing the ratio from Bottom-25% to Bottom-50% drops the General Average from **61.34** to **57.93**, without improving safety. This suggests that the Top-25% of Low-Conflict heads are sufficient for alignment, and adding more simply introduces noise.
- **The “Full SFT” Cliff:** Interestingly, within the Risky Zone, performance recovers slightly as we expand from Top-25% to Top-75% (MMLU rises to 50.98), likely because more low-conflict heads

are included. However, this trend ends abruptly at 100% (Full SFT), where the General Average crashes to **36.78**. This sharp drop confirms that blindly updating every parameter triggers the worst-case conflict, far exceeding the harm of updating just the risky subset.

E.3 Sensitivity to Data Size: Data Efficiency

To assess whether our conflict diagnosis requires large-scale data, we restrict the calibration set to only **100 samples** from MMLU (a $5\times$ reduction). We evaluate this setting on both Llama and Qwen to verify cross-model consistency.

Results. As detailed in Table 10, despite the minimal data, the conflict metric remains highly effective in identifying the optimal utility zone.

- **Llama:** We observe a strictly monotonic trend. The General Average improves progressively from 55.61% in the Risky Zone (\mathcal{B}_1) to **63.18%** in the Safe Zone (\mathcal{B}_4). Notably, \mathcal{B}_4 achieves the best performance across all individual utility tasks (MMLU, CSQA, GSM8K, MATH).
- **Qwen:** Similarly, the Low-Conflict bucket (\mathcal{B}_4) yields the highest General Average of **74.71%**, significantly outperforming the intermediate buckets (\mathcal{B}_2 : 64.94%).

This confirms that the "Conflict Landscape" for general knowledge is a salient structural property that can be accurately captured even with extremely sparse signals, ensuring the method’s efficiency and reproducibility across different architectures.

E.4 Sensitivity to Data Source: The “Inverted U-Shape” Phenomenon

While MMLU calibration yields consistent trends favoring the "Low Conflict" tail, we further investigate whether these conflict maps transfer across domains. We replace the calibration source with **100 samples from GSM8K** to diagnose conflict specifically for reasoning tasks.

Results. As presented in Table 11, switching the data source reveals a distinct “**Inverted U-shape**” pattern, where the optimal trade-off shifts from the tail (\mathcal{B}_4) to intermediate buckets.

- **Llama:** The utility peak appears at \mathcal{B}_3 (**62.05%** Gen Avg), with GSM8K accuracy

Sparsity Ratio	General Utility				Safety Performance			Average	
	MMLU	CSQA	GSM8K	MATH	AdvHarm	WildGuard	DAN	Gen Avg	Safe Avg
Full SFT (100%)	46.28	60.84	20.80	19.20	93.26	94.29	84.27	36.78	90.61
<i>Risky Zone Trajectory (Prioritizing High-Conflict Heads)</i>									
Top-25%	48.52	63.83	68.80	42.93	95.22	96.64	83.51	56.02	91.79
Top-50%	47.79	58.72	70.60	42.00	96.30	97.61	89.54	54.78	94.48
Top-75%	50.98	67.16	73.20	40.00	96.75	97.35	90.11	57.83	94.73
<i>Safe Zone Trajectory (CAST, Prioritizing Low-Conflict Heads)</i>									
Bottom-25% (Ours)	55.73	69.42	77.20	43.00	94.85	96.69	86.33	61.34	92.62
Bottom-50%	53.67	67.65	68.00	42.40	94.90	95.23	86.33	57.93	92.15
Bottom-75%	53.24	65.85	61.80	43.00	96.70	97.61	90.89	55.97	95.07

Table 9: **Sensitivity Analysis to Sparsity Ratio on Llama.** We compare the performance trajectory when prioritizing High-Conflict heads (Top-k%) vs. Low-Conflict heads (Bottom-k%). The Bottom-25% setting achieves the optimal trade-off, maximizing general utility while maintaining robust safety.

Model	Bucket	General Capabilities				Avg.	Safety Performance			Avg.
		MMLU	CSQA	GSM8K	MATH	General	AdvHarm	WildGuard	DAN	Safety
Llama (MMLU-100)	\mathcal{B}_1 (High C)	50.61	64.62	66.00	41.20	55.61	96.45	97.88	88.90	94.41
	\mathcal{B}_2	50.63	67.32	71.60	42.20	57.94	97.15	98.28	92.88	96.10
	\mathcal{B}_3	54.27	64.95	73.80	43.20	59.05	95.15	96.42	80.07	90.55
	\mathcal{B}_4 (Low C)	56.89	70.43	81.00	44.40	63.18	94.10	97.35	85.20	92.21
Qwen (MMLU-100)	\mathcal{B}_1 (High C)	68.00	81.16	86.80	59.60	73.89	83.80	89.12	81.00	84.64
	\mathcal{B}_2	68.09	81.08	58.80	51.80	64.94	88.60	89.79	81.07	86.49
	\mathcal{B}_3	68.37	81.16	80.80	60.60	72.73	69.30	82.89	78.58	76.92
	\mathcal{B}_4 (Low C)	68.47	81.57	80.40	68.40	74.71	67.70	82.76	76.94	75.80

Table 10: **Data Efficiency Analysis (MMLU-100 Calibration).** We report full metrics for conflict diagnosis using only 100 MMLU samples. Across both Llama and Qwen, the Low-Conflict bucket (\mathcal{B}_4) consistently identifies the optimal region for General Utility, validating the robustness of the metric under sparse data conditions.

Model	Bucket	General Capabilities				Avg.	Safety Performance			Avg.
		MMLU	CSQA	GSM8K	MATH	General	AdvHarm	WildGuard	DAN	Safety
Llama (GSM8K-100)	\mathcal{B}_1 (High C)	56.52	69.37	70.20	41.60	59.42	96.40	97.08	86.62	93.37
	\mathcal{B}_2	53.64	65.03	72.60	42.40	58.42	96.25	96.68	93.74	95.56
	\mathcal{B}_3 (Optimal)	54.66	68.55	83.40	41.60	62.05	96.55	98.54	89.75	94.95
	\mathcal{B}_4 (Low C)	51.72	67.32	63.40	36.40	54.71	95.85	96.29	87.83	93.32
Qwen (GSM8K-100)	\mathcal{B}_1 (High C)	68.26	81.74	75.40	66.80	73.05	77.40	88.20	74.16	79.92
	\mathcal{B}_2 (Optimal)	68.94	81.08	86.60	67.40	76.00	82.15	88.59	78.22	82.99
	\mathcal{B}_3	67.36	81.82	44.60	51.00	61.19	86.30	89.79	78.15	84.75
	\mathcal{B}_4 (Low C)	67.74	80.51	80.20	61.80	72.56	73.80	86.34	74.66	78.27

Table 11: **Domain Transfer Analysis (GSM8K-100 Calibration).** We switch the diagnostic source to 100 samples from GSM8K. Unlike the monotonic trend in MMLU calibration, both Llama and Qwen exhibit an “Inverted U-shape” or intermediate peak (\mathcal{B}_3 and \mathcal{B}_2 , respectively), confirming that reasoning tasks require a specific trade-off zone distinct from the extreme tail.

reaching **83.40%**. This significantly outperforms the Low-Conflict tail \mathcal{B}_4 (GSM8K: 63.40%, Gen Avg: 54.71%).

- **Qwen:** A similar structural "Sweet Spot" is observed at \mathcal{B}_2 . It achieves the highest General Average (**76.00%**) and GSM8K score (**86.60%**), surpassing the \mathcal{B}_4 bucket (72.56% and 80.20% respectively).

Analysis: The Limitation of Single-Head Ablation. We attribute this “Inverted U-shape” to the

limitations of using single-head ablation to measure sensitivity in reasoning tasks. Unlike factual queries, reasoning relies on complex, multi-step logical chains. Masking a single head often fails to break this chain due to model redundancy, leading our metric $\mathcal{S}(h)$ to underestimate the head’s importance. Consequently, many critical reasoning heads are incorrectly assigned low scores and grouped into the “Low Conflict” bucket (\mathcal{B}_4). However, when these heads are modified during fine-tuning, the subtle changes disrupt the precise logic flow.

Model	Zone Type	MMLU	CSQA	GSM8K	MATH	Adv Harm	WildG	DAN	Avg. Gen	Avg. Safe
Llama-3.1 (8B)	Risky Zone	19.43	11.06	12.00	33.40	95.90	96.95	93.45	18.97	95.43
	Safe Zone	50.82	57.82	42.00	31.80	95.50	96.95	86.55	45.61	93.00
	FULL SFT	17.10	17.20	40.00	23.20	96.55	97.35	97.44	24.37	97.11
Qwen2.5 (7B)	Risky Zone	66.76	80.75	42.00	50.80	90.70	94.43	92.53	60.08	92.55
	Safe Zone	66.69	81.16	68.20	56.60	91.80	94.16	91.81	68.16	92.59
	FULL SFT	66.95	78.95	62.20	55.80	90.40	96.15	88.90	65.98	91.82
Mistral (7B)	Risky Zone	35.26	43.41	9.20	4.00	76.00	88.73	69.04	22.97	77.92
	Safe Zone	38.40	44.31	11.80	2.80	74.20	88.46	63.20	24.33	75.29
	FULL SFT	45.51	54.05	22.40	9.20	46.40	72.41	66.55	32.79	61.79

Table 12: Results for Tuning $W_k + W_v$ Projections.

Model	Method	MMLU	CSQA	GSM8K	MATH	Adv Harm	WildG	DAN	Avg. Gen	Avg. Safe
Llama	Base	59.38	73.23	87.39	44.40	55.64	81.60	64.40	66.10	67.22
	FULL-DPO	58.31	71.25	63.20	44.80	78.20	89.66	74.52	59.39	80.79
	CAST-DPO (Safe)	58.33	71.25	82.20	42.00	92.60	93.77	70.89	63.44	85.75
Qwen	Base	68.63	80.50	91.00	67.40	34.74	70.41	66.83	76.88	57.33
	FULL-DPO	68.03	80.02	89.80	69.40	46.60	76.66	70.96	76.81	64.74
	CAST-DPO (Safe)	68.29	80.59	90.40	70.00	52.90	80.77	73.95	77.32	69.21
Mistral	Base	46.98	59.13	27.80	8.00	32.79	64.45	53.10	35.48	50.11
	FULL-DPO	45.76	56.84	26.00	7.80	40.70	68.44	55.87	34.10	55.00
	CAST-DPO (Safe)	45.79	56.35	18.80	7.40	54.35	77.59	63.49	32.08	65.14

Table 13: DPO Generalization Results on PKU-SafeRLHF Dataset.

Qwen3-30B	MMLU	CSQA	GSM8K	MATH	Adv Harm	Wild Guard	DAN	Avg. Gen	Avg. Safe
Base	69.79	79.44	87.20	56.00	78.86	83.02	71.75	73.11	77.88
FULL-SFT	68.33	77.23	77.80	44.40	84.60	88.86	86.34	66.94	86.60
CAST-SFT (Risky)	68.65	78.13	79.40	48.20	84.20	87.53	84.70	68.60	85.48
CAST-SFT (Safe)	69.36	79.11	84.20	54.00	82.40	86.73	85.41	71.67	84.85

Table 14: Scalability Results on Qwen3-30B-A3B-Instruct (MoE).

As a result, the true “safe zone” shifts to intermediate buckets, where the balance between safety and valid sensitivity estimation is better maintained.

E.5 Generalization to W_k and W_v Projections

To address the concern regarding parameter-type coverage, we extended CAST to Key (W_k) and Value (W_v) projections. For GQA-based models, we adapted our diagnosis by aggregating $C(h)$ scores from all W_q heads within a GQA group to profile their shared $W_{k,v}$ component. As shown in Table 12, tuning the Safe Zone in $W_{k,v}$ consistently outperforms Full SFT, effectively preserving utility across Llama, Qwen, and Mistral families.

E.6 Generalization to Preference Optimization (DPO)

We further validated CAST using the DPO objective on the PKU-SafeRLHF dataset (Ji et al., 2024). Results in Table 13 demonstrate that Safe-Zone

DPO maintains reasoning performance (GSM8K) significantly more effectively than Full DPO while achieving comparable safety win-rates. This confirms that structural conflict patterns identified by CAST are objective-invariant.

E.7 Scalability to Large-Scale MoE Models (30B)

We validated CAST on the **Qwen3-30B-A3B-Instruct** MoE model to test its scalability. As shown in Table 14, CAST-SFT (Safe Zone) achieved an Avg. Gen score of **71.67**, significantly outperforming Full SFT (66.94) and the Risky Zone (68.60). This proves CAST’s precision in identifying high-conflict parameters within sparse, high-parameter modular architectures.

E.8 Discussion on MLP Components

While this study primarily focuses on attention heads, we acknowledge that MLPs are critical

for factual storage. Preliminary results on applying CAST to MLP down-projections show consistent conflict localization, suggesting that the principle of modular heterogeneity extends to feed-forward components. We will incorporate a formal MLP-extension analysis in the final version of the manuscript.

E.9 Computational Complexity

Benchmarked on Llama (NVIDIA H800), analyzing a single head (100 samples) takes approximately 7 seconds for utility (next-token prediction) and 1 minute for safety (generation). For a model with N heads, the total time scales linearly to $N \times \text{Cost}_{\text{head}}$. While this represents a fixed upfront cost, it is a strict **one-off pre-computation**. Unlike random search strategies that require expensive, repeated training runs to stumble upon an effective sparsity pattern, our approach generates a precise Conflict Map in a single pass. This makes CAST not only more interpretable but also significantly more cost-effective than blind trial-and-error.