A Geometric Perspective on Stabilizing Value Conflict Resolution

Saket Reddy SAKETR3@ILLINOIS.EDU
Supervised Program for Alignment Research (SPAR), University of Illinois - Urbana-Champaign

Andy Liu ANDYLIU@ANDREW.CMU.EDU
Carnegie Mellon University

Abstract

Large Language Models (LLMs) often struggle to navigate user requests that contain complex value conflicts. To address this challenge, we investigate how chain-of-thought (CoT) reasoning can help improve performance in this domain. Geometrically, we show that CoT correlates with further smoothing the model's loss landscape in its sharpest direction, helping resolve the optimization instability of traditional scalar rewards. In terms of performance, we demonstrate via relevant downstream benchmarks that CoT may increase capability, demonstrating that value conflict-focused CoT has the potential to be an effective mechanism for better moral reasoning. To capitalize on this potential, we create a new value conflict-focused CoT design that further smooths the sharpest direction of the loss landscape and increases moral reasoning performance. This finding shows that explicitly modifying and improving the design of reasoning dynamics offers a promising avenue for advancing pluralistic alignment in LLMs. [1]

1. Introduction

User requests often require navigating diverse and conflicting values, such as balancing helpfulness with harmlessness [1]. When models are forced to optimize these contradictory constraints simultaneously through scalar rewards in Reinforcement Learning from Human Feedback (RLHF), model behavior can become unstable [2, 20]. In this work, we investigate how chain-of-thought (CoT) can help resolve competing values. We first apply a geometric analysis, where we find that CoT correlates with further smoothing the model's loss landscape beyond the effects of supervised fine-tuning (SFT), helping mitigate the instability caused by direct scalar reward training. We then apply a performance analysis, where we find that CoT may increase capability on downstream benchmarks. This performance increase is within the margin of error, but the trend is consistent, showing that value conflict-focused CoT has the potential to be an effective tool for better moral reasoning.

We then further explore the effects of modifying the design of reasoning dynamics. Value conflict resolution in Large Language Models (LLMs) can be thought of similarly to geometric frustration: a physics phenomenon that occurs when a system's constraints cannot be satisfied simultaneously, leading to a spiky energy landscape [12]. When a scalar reward model forces a policy to optimize for mutually exclusive values without a resolution mechanism, the model can be thought of as experiencing similar frustration. Geometric frustration is often resolved via annealing: a thermodynamic process of heating the system to a high temperature, followed by gradual cooling to freeze the system into a stable configuration [13, 16]. We loosely translate this thermodynamic principle into an annealing-inspired CoT framework that enforces a cognitive transition from high-

1. We release our dataset and code.

© S. Reddy & A. Liu.

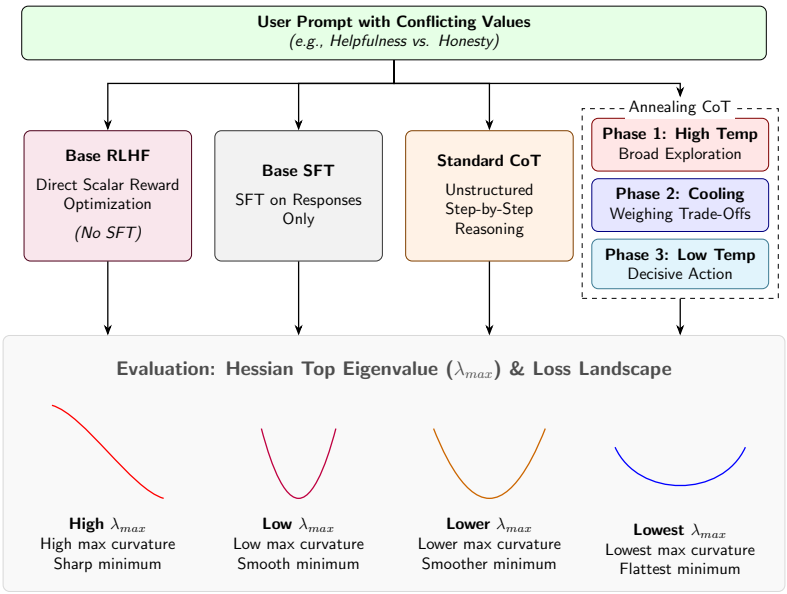

Figure 1: Overview of our training methods and their corresponding effects on loss landscape geometry. Base RLHF leads to unstable value conflict resolution, indicated by a high top Hessian eigenvalue and high curvature of the loss landscape. Base SFT smooths the landscape, indicated by a low top eigenvalue and low maximum curvature. Both standard and annealing CoT lead to even more smooth landscapes, indicated by an even lower top eigenvalue and maximum curvature.

temperature exploration (broadly weighing values) to low-temperature convergence (making decisive, hierarchical trade-offs). We find that this modified CoT design corresponds to a flatter minimum in the loss landscape and achieves better performance by a larger margin. This shows that more intentionally designing reasoning dynamics offers an encouraging direction for advancing pluralistic alignment.

Our contributions are as follows:

- We demonstrate that resolving value conflicts directly via RLHF creates an unstable loss landscape.
- We investigate how CoT resolves this instability. Geometrically, we find that CoT further smooths the loss landscape beyond the effects of SFT. In terms of capability, we observe that CoT may increase moral reasoning performance, though this increase is within the margin of error. This demonstrates that value conflict-focused CoT has the potential to be a good mechanism for better moral reasoning.
- We demonstrate that a new value conflict-focused CoT design corresponds to a flatter minimum in the loss landscape and achieves better performance by a larger margin, indicating the importance of more intentional CoT design.

We release our standard and annealing CoT dataset. An overview of our training methods and their corresponding effects on the model's loss landscape geometry is shown in Figure 1.

2. Background and Related Work

2.1. CoT for Alignment

RLHF often struggles to navigate complex objectives due to reliance on scalar rewards. To address this, there have been efforts to develop specialized CoT frameworks for alignment [6, 26]. For instance, Jin et al. [10] introduce a framework to chain together multiple prompts to elicit structured ethical deliberations. However, while these approaches successfully improve general moral reasoning, the explicit design of new CoT variants for resolving direct value conflicts specifically remains underexplored. This work addresses this gap by exploring how intentionally modifying reasoning dynamics can help models better navigate competing values, and investigating the underlying geometry behind this.

2.2. Loss Landscapes and Hessian Analysis

Loss landscape analysis is a fundamental method for understanding optimization dynamics. In the context of LLMs, Chen et al. [3] show that effective post-training molds the landscape into smooth, basin-like structures. Following Ghorbani et al. [5], we also analyze the eigenvalues of the loss function's Hessian matrix, $\nabla^2 L(\theta)$, to quantitatively assess model stability. Because the Hessian captures second-order partial derivatives, it measures how sharply the loss landscape bends across parameter directions. To evaluate this curvature, we examine the Hessian's top eigenvalue ($\lambda_{\max}$) [11]. The top eigenvalue represents the maximum curvature in the landscape's sharpest, most unstable direction. If $\lambda_{\max}$ is small, then small weight updates will not drastically increase the loss, indicating stability. Geometrically, the landscape is wide and gently sloping, similar to a basin. Conversely, a large $\lambda_{\max}$ means that small steps will cause the loss to spike, indicating a steep, unstable, cliff-like minimum. Because a lower $\lambda_{\max}$ mathematically indicates a flatter, more stable minimum, analyzing the Hessian is a well-established technique for evaluating training dynamics [9, 28].

3. Methodology

3.1. Data Generation

We start by utilizing prompts generated by the ConflictScope pipeline [15], an automated pipeline designed to elicit value conflicts. Specifically, we use the authors' publicly released prompt sets, which were constructed by prompting Claude 3.5 Sonnet to generate conflicting scenarios for diverse value pairs. Using an ensemble of strong LLMs (GPT 5 Nano, Gemini 3.1 Flash Lite, and Claude 4.5 Haiku), we synthetically constructed two reasoning traces and corresponding responses for each prompt to perform SFT on. The first generated CoT-response pair for each prompt is a standard CoT and response, where the model resolves the value conflict through unstructured reasoning. The second generated CoT-response pair for each prompt is an annealing-inspired CoT and response, where the reasoning trace is structured as follows:

- *Phase 1 (High Temperature):* A broad exploration state where the model explicitly weighs conflicting constraints without prematurely collapsing the solution space.
- *Phase 2 (Cooling):* A transitional state that evaluates trade-offs and applies hierarchical principles to systematically narrow the available actions.

- *Phase 3 (Low Temperature):* A convergence state that finalizes a stable, decisive action based on the narrowed parameters.

We also gave the LLMs length requirements for their generation to control for length. Examples of the CoTs for SFT are shown in Appendix H, and the prompts used to generate the CoTs are shown in Appendix G. To create the annealing-inspired CoTs, the text from each of the three phases was appended together.

3.2. Training

We utilize Qwen 3.5 (0.8B) [18], an instruct-tuned model, for fine-tuning. Due to compute constraints, we utilize this sub-billion parameter model to ensure the heavy Hessian computations were feasible. We also intentionally utilized a full fine-tune instead of parameter-efficient fine-tuning (PEFT) so that the true loss landscapes were not artificially altered. We follow a standard post-training pipeline. For our CoT models, we first perform SFT on the prompt, CoT, and response triplets, then perform RLHF using Group Relative Policy Optimization (GRPO) [19]. Our training hyperparameters are detailed in Appendix D. We employ GRM Llama 3.2 (3B) to provide scalar rewards [23].

We also train two alternate baseline models to isolate the gains from CoT. The first alternate model is a base SFT model, which undergoes SFT on responses only without CoT, then undergoes RLHF. The second alternate model is a base RLHF model, which undergoes RLHF only on the prompts from both the SFT and RLHF phases of the CoT training, without any SFT.

3.3. Evaluation

To quantify the stability of the local minima achieved by each method, we analyze the maximum curvature of the loss landscape on a held-out evaluation subset of ConflictScope. Because exact Hessian computation is intractable for large networks, we utilize PyHessian [24] to compute stochastic approximations of the eigenvalues via Hutchinson's method [8]. To ensure statistical robustness, we compute the top eigenvalue ($\lambda_{\max}$) across three distinct subsets of the evaluation data, reporting the aggregated means and standard error. We also construct loss landscapes for visualization. To rigorously capture the true structural stability of the local minima, we explicitly perturb the trained model weights along the top two eigenvectors extracted from the Hessian matrix, following the foundational methodology established by Yao et al. [24]. More details on how the loss landscapes were created can be found in Appendix F.

We also evaluate the models on three performance benchmarks (setup details in Appendix E):

- **MMLU Moral Scenarios** [7]: Evaluates whether models classify different simple actions as moral or immoral.
- **DecodingTrust Machine Ethics** [21]: Utilizing the commonsense morality subset, this benchmark evaluates how models navigate ethical dilemmas. Unlike MMLU Moral Scenarios, this dataset also includes complex real-world situations sourced from Reddit, providing an assessment of model performance under nuanced, real-life value conflicts.
- **SafetyBench** [27]: Evaluates model boundaries by testing opinions on sensitive topics (e.g. mental health advice, privacy protection, etc.), ensuring the model answers safely even for sensitive questions.

4. Results

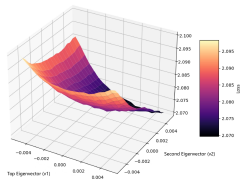

(a) Base RLHF

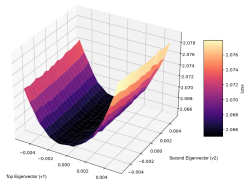

(b) Base SFT

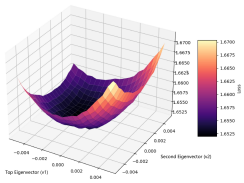

(c) Standard CoT

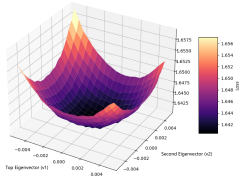

(d) Annealing CoT

Figure 2: 3D eigenplane loss landscapes generated by perturbing model weights along the top two eigenvectors (v_1, v_2). The base RLHF model displays an unstable cliff-like landscape in its sharpest direction. The base SFT model displays a smooth ravine-like landscape. Finally, the standard CoT and annealing CoT models display more smooth, basin-like landscapes in their sharpest directions.

Table 1: Top Hessian eigenvalues (lower eigenvalues are better); $\pm$ denotes one standard error. The base RLHF model has a high top eigenvalue indicating sharpness (instability), while the base SFT model has a low top eigenvalue indicating flatness (stability). The standard CoT model has a lower top eigenvalue, and the annealing CoT model has an even lower top eigenvalue, indicating an even flatter minimum.

Model	**Top Eigenvalue (λ_{max})**
Base RLHF	4083.87 ± 156.53
Base SFT	1402.77 ± 4.48
Standard CoT	1345.65 ± 15.60
Annealing CoT	1218.04 ± 43.31

The loss landscape visualizations are shown in Figure 2 and the quantitative eigenvalue results are shown in Table 1. 2D contour maps of the visualizations are shown in Appendix A. When the base RLHF-only model is forced to resolve conflicting values directly via scalar reward optimization, it has a high top eigenvalue ($\lambda_{\text{max}} \approx 4083$) and cliff-like loss landscape (Figure 2a), demonstrating instability. In contrast, by performing SFT on responses only, the top eigenvalue of the base SFT-trained model drops to $\lambda_{\text{max}} \approx 1402$ and the loss landscape erodes from a loss "cliff" into a smoother ravine-like landscape (Figure 2b).

While SFT effectively mitigates instability, CoT correlates with better geometric convergence. The standard CoT model achieves a lower top eigenvalue ($\lambda_{\text{max}} \approx 1345$), indicating that reasoning

Table 2: Benchmark performance across 0.8B models. Results represent the mean accuracy $\pm$ one standard error. The base RLHF model achieves the lowest performance, the base SFT and standard CoT models achieve better performance within the margin of error, and the annealing CoT model achieves the highest performance by a larger margin, indicating the importance of CoT design for moral reasoning capability.

Model	**MMLU Moral Scenarios**	**DecodingTrust Machine Ethics**	**SafetyBench**
Base RLHF	24.67 ± 0.88	42.33 ± 1.20	43.67 ± 4.37
Base SFT	29.67 ± 0.33	43.67 ± 0.33	48.67 ± 1.76
Standard CoT	30.33 ± 2.40	45.67 ± 2.03	53.67 ± 5.46
Annealing CoT	$\mathbf{41.67 \pm 0.88}$	$\mathbf{54.33 \pm 2.85}$	$\mathbf{64.00 \pm 2.08}$

flattened the sharpest curvature further. Visually, this corresponds to a wide and gently sloping maximum curvature in the landscape's sharpest direction, as shown by the smooth, stable basin in Figure 2c. Finally, the annealing CoT model demonstrates even better geometric stability. While there is no significant visual difference between the loss landscapes of the standard CoT and annealing CoT models (Figure 2d), the annealing CoT achieves a lower top eigenvalue ($\lambda_{\text{max}} \approx 1218$), indicating that the annealing structure successfully flattened the dominant direction of curvature even further.

We also ran benchmark evaluations. The results are shown in Table 2. The base RLHF model, which exhibited the most unstable minimum in its sharpest direction, obtained the lowest performance across all three benchmarks. The base SFT model, which smoothed the loss landscape in its sharpest direction, demonstrated consistent improvements, though these improvements are within the margin of error. The standard CoT model further improved on the performance of the base SFT model. However, its improvement is again within the margin of error, suggesting that while CoT has the potential to improve capability, it is not certain. Finally, the annealing CoT model, which correlated with the flattest, most stable minimum in the model's sharpest direction, outperformed both baselines by a larger margin. Increasing performance on these benchmarks while only fine-tuning on value conflict-related CoTs suggests that, when designed more optimally, value conflict-focused CoT does improve the model's ability to navigate moral reasoning. Ablation studies on an alternate model family and larger model scale are shown in Appendices B and C, respectively.

5. Conclusion

LLMs often struggle when they encounter value conflicts. We investigated CoT as a mechanism to better resolve these competing tensions. Geometrically, our analysis demonstrated that while scalar rewards induce high optimization instability, CoT correlates with a lower top Hessian eigenvalue and smoothing the loss landscape in its sharpest direction beyond the effects of SFT into a stable basin. In terms of performance, we observe that value conflict-focused CoT has the potential to improve moral reasoning performance. Furthermore, we demonstrated that explicitly modifying reasoning dynamics yields even better performance. By introducing an annealing CoT, the model found an even flatter, more stable minimum, indicated by a lower top Hessian eigenvalue. The annealing CoT model also improved moral reasoning performance by a larger margin. This finding highlights that more intentionally designing the internal reasoning dynamics of models offers a promising avenue for fulfilling more robust pluralistic alignment.

References

[1] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

[2] Han Bao, Yue Huang, Xiaoda Wang, Zheyuan Zhang, Yujun Zhou, Carl Yang, Xiangliang Zhang, and Yanfang Ye. Ai alignment breaks at the edge. *arXiv preprint arXiv:2602.20042*, 2026.

[3] Huanran Chen, Yinpeng Dong, Zeming Wei, Yao Huang, Yichi Zhang, Hang Su, and Jun Zhu. Unveiling the basin-like loss landscape in large language models. *arXiv preprint arXiv:2505.17646*, 2025.

[4] Jiaqing Chen, Nicholas Hadler, Tiankai Xie, Rostyslav Hnatyshyn, Caleb Geniesse, Yaoqing Yang, Michael Mahoney, Talita Perciano, John Hartwig, Ross Maciejewski, and Gunther Weber. Landscaper: Understanding loss landscapes through multi-dimensional topological analysis. *arXiv preprint arXiv:2602.07135*, 2026.

[5] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2232–2241, 2019.

[6] Melody Y. Guan, Manas Joglekar, Eric Wallace, Saachi Jain, Boaz Barak, Alec Helyar, Rachel Dias, Andrea Vallone, Hongyu Ren, Jason Wei, Hyung Won Chung, Sam Toyer, Johannes Heidecke, Alex Beutel, and Amelia Glaese. Deliberative alignment: Reasoning enables safer language models. *arXiv preprint arXiv:2412.16339*, 2024.

[7] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[8] Michael F. Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 19(2):433–450, 1990.

[9] Ajay Jaiswal, Yifan Wang, Lu Yin, Shiwei Liu, Runjin Chen, Jiawei Zhao, Ananth Grama, Yuandong Tian, and Zhangyang Wang. From low rank gradient subspace stabilization to low-rank weights: Observations, theories, and applications. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.

[10] Zhijing Jin, Sydney Levine, Fernando Gonzalez, Ojasv Kamal, Maarten Sap, Mrinmaya Sachan, Rada Mihalcea, Josh Tenenbaum, and Bernhard Scholkopf. When to make exceptions: Exploring language models as accounts of human moral judgment. In *Advances in Neural Information Processing Systems*, volume 35, pages 27376–27390, 2022.

[11] Anil Kamber and Rahul Parhi. Sharpness of minima in deep matrix factorization: Exact expressions. *arXiv preprint arXiv:2509.25783*, 2025.

[12] Woo Jin Kim, Michelle A. Smeaton, Chunjing Jia, Berit H. Goodge, Byeong-Gwan Cho, Kyuho Lee, Motoki Osada, Daniel Jost, Anton V. Ievlev, Brian Moritz, Lena Kourkoutis, Thomas Devereaux, and Harold Y. Hwang. Geometric frustration of jahn-teller order in the infinite-layer lattice. *Nature*, 615:237–243, 2023.

[13] Scott Kirkpatrick, C. Daniel Gelatt, and Mario P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

[14] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[15] Andy Liu, Kshitish Ghate, Mona Diab, Daniel Fried, Atoosa Kasirzadeh, and Max Kleiman-Weiner. Generative value conflicts reveal llm priorities. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026.

[16] Yutong Luo, Yi-Zheng Zhen, Xiangjing Liu, Daniel Ebler, and Oscar Dahlsten. Bound on annealing performance from stochastic thermodynamics, with application to simulated annealing. *Physical Review E*, 108(5):054119, 2023.

[17] Meta. Llama-3.2-1B, 2024.

[18] Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026.

[19] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

[20] Ying Su, Mingen Zheng, Weili Diao, and Haoran Li. Jailbreaking large language models with morality attacks. *arXiv preprint arXiv:2604.17053*, 2026.

[21] Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, Sang T. Truong, Simran Arora, Mantas Mazeika, Dan Hendrycks, Zinan Lin, Yu Cheng, Sanmi Koyejo, Dawn Song, and Bo Li. DecodingTrust: A comprehensive assessment of trustworthiness in GPT models. In *Advances in Neural Information Processing Systems (Datasets and Benchmarks Track)*, volume 36, pages 73188–73248, 2023.

[22] Tiankai Xie, Caleb Geniesse, Jiaqing Chen, Yaoqing Yang, Dmitriy Morozov, Michael W. Mahoney, Ross Maciejewski, and Gunther H. Weber. Evaluating loss landscapes from a topology perspective. In *Advances in Neural Information Processing Systems*, 2024.

[23] Rui Yang, Ruomeng Ding, Yong Lin, Huan Zhang, and Tong Zhang. Regularizing hidden states enables learning generalizable reward model for llms. In *Advances in Neural Information Processing Systems*, 2024.

[24] Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W. Mahoney. Pyhessian: Neural networks through the lens of the hessian. *arXiv preprint arXiv:1912.07145*, 2019.

[25] Youzhi Yu. llama3.2-1b-thinking, 2025. URL https://huggingface.co/PursuitOfDataScience/llama3.2-1b-thinking.

[26] Yichi Zhang, Siyuan Zhang, Yao Huang, Zeyu Xia, Zhengwei Fang, Xiao Yang, Ranjie Duan, Dong Yan, Yinpeng Dong, and Jun Zhu. STAIR: Improving safety alignment with introspective reasoning. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.

[27] Zhexin Zhang, Leqi Lei, Lindong Wu, Rui Sun, Yongkang Huang, Chong Long, Xiao Liu, Xuanyu Lei, Jie Tang, and Minlie Huang. SafetyBench: Evaluating the safety of large language models with multiple choice questions. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15582–15621. Association for Computational Linguistics, 2024.

[28] Shuchen Zhu, Rizhen Hu, Mingze Wang, Mou Sun, Xue Wang, Kun Yuan, and Zaiwen Wen. Accelerating llm pre-training through flat-direction dynamics enhancement. *arXiv preprint arXiv:2602.22681*, 2026.

Appendix A. Main Study 2D Contour Maps

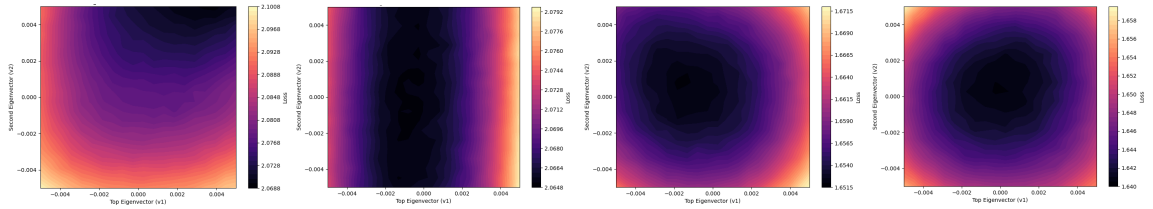

(a) Base RLHF (b) Base SFT (c) Standard CoT (d) Annealing CoT

Figure 3: 2D contour maps of the eigenplane loss landscapes generated by perturbing model weights along the top two eigenvectors (v_1, v_2). The map for the base RLHF model depicts an incline, indicating instability. The map for the base SFT model depicts a ravine, indicating more stability. Finally, the maps for the standard and annealing CoT models depict a circular valley, indicating further stability.

Appendix B. Alternate Model Family Ablation Study

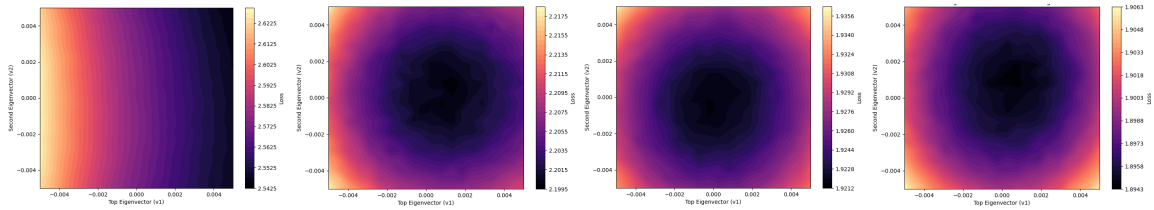

(a) Base RLHF (b) Base SFT (c) Standard CoT (d) Annealing CoT

Figure 4: 2D contour maps of the eigenplane loss landscapes for the Llama 3.2 (1B) models generated by perturbing model weights along the top two eigenvectors (v_1, v_2). Similarly to the contour maps for Qwen3.5, the map for the base RLHF model depicts an incline (indicating instability), while the map for the standard and annealing CoT models depict circular valleys (indicating stability). Unlike the map for Qwen 3.5, the base SFT model also depicts a stable circular valley, although the top Hessian eigenvalues for the CoT models are still lower than for the base SFT model, indicating more stability.

To ensure our findings generalize to another model family, we performed an ablation study where we replicated our pipeline on Llama 3.2 (1B) [17]. We again used a small, billion parameter model to keep the Hessian computations feasible. Additionally, as Llama 3.2 is not a native reasoning model which thinks using CoT, we used a version of Llama 3.2 that was fine-tuned to add CoT reasoning capabilities [25], ensuring a fairer comparison with the Qwen 3.5 model used in the original study. Similarly to Qwen 3.5 0.8B, this model does not usually think in reasoning streams before answering non-math and coding questions, allowing for comparison of non-reasoning and reasoning performance on the same model architecture. The loss landscape visualizations are shown in Figures 5 and 4, the quantitative eigenvalue results are shown in Table 3, and the benchmarks results are shown in Table 4.

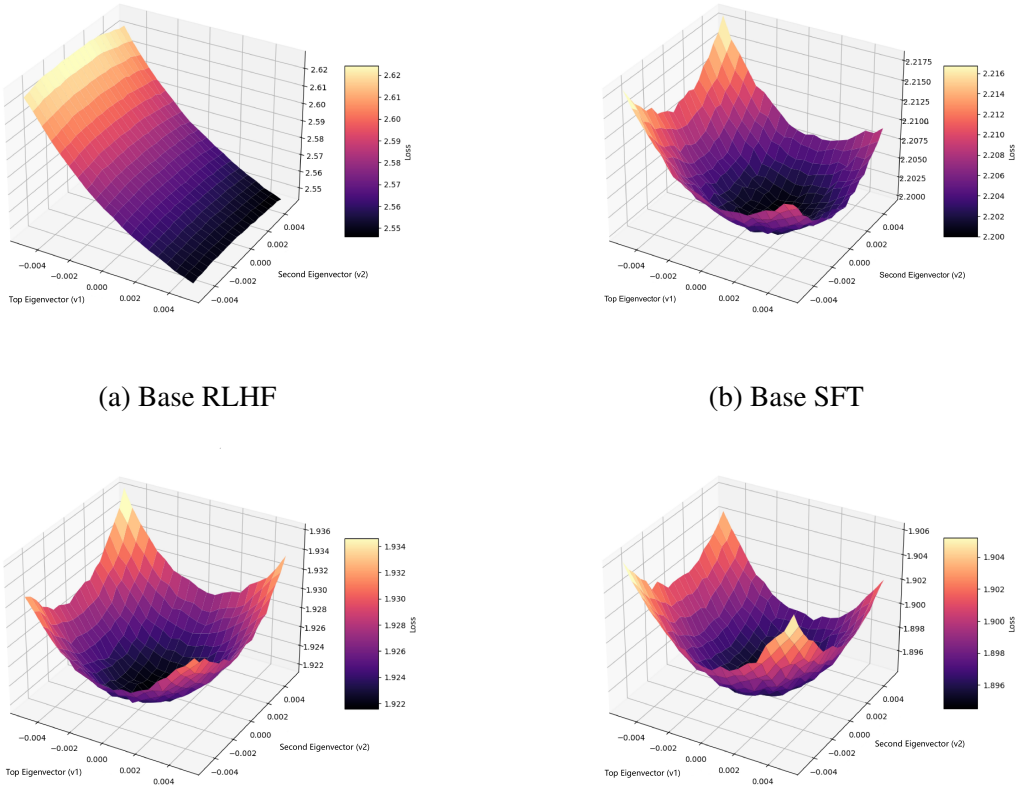

(a) Base RLHF

(b) Base SFT

(c) Standard CoT

(d) Annealing CoT

Figure 5: 3D eigenplane loss landscapes for Llama 3.2 (1B) generated by perturbing model weights along the top two eigenvectors (v_1, v_2). The base RLHF model displays a cliff-like minimum (indicating instability), while the standard and annealing CoT models display more basin-like minima (indicating stability). Unlike Qwen 3.5, the base SFT model also displays a stable basin, although the CoT models still achieve lower top Hessian eigenvalues, indicating more stability.

Geometrically, we found that the observations we made on Qwen3.5 remain similar for this alternate model family. The base model exhibited an unstable cliff-like landscape in its sharpest direction with a high top eigenvalue, while the standard CoT model exhibited a stable basin-like landscape with a lower top eigenvalue. The annealing CoT model further lowered the top eigenvalue, aligning with an even flatter, more stable minimum in the sharpest direction of the landscape. The only notable difference is that the base SFT model also aligns with a more stable basin-like landscape instead of the ravine-like landscape of Qwen 3.5. However, the standard and annealing CoT models achieve lower top eigenvalues than the base SFT model, indicating that CoT still further smooths the loss landscape.

Regarding performance, the observations remain the same as for Qwen 3.5. The base RLHF model obtains the lowest performance while the base SFT and standard CoT models demonstrate improvements within the margin of error (although standard CoT does help to a large extent on the

Table 3: Top Hessian eigenvalues for Llama 3.2 1B (lower eigenvalues are better); $\pm$ denotes one standard error. The base RLHF model has a high top eigenvalue (indicating instability), while the base SFT model has a low top eigenvalue (indicating stability). The standard CoT model further reduces top eigenvalue, and the annealing CoT model has an even lower top eigenvalue, indicating an even flatter minimum.

Model	**Top Eigenvalue (λ_{max})**
Base RLHF	3012.05 ± 64.08
Base SFT	951.87 ± 6.53
Standard CoT	842.51 ± 3.68
Annealing CoT	809.36 ± 1.46

Table 4: Benchmark performance across Llama 3.2 (1B) models. Results represent the mean accuracy $\pm$ one standard error. Consistent with earlier results, the base RLHF models achieve the lowest performance. The base SFT and standard CoT models achieve better performance within the margin of error. Finally, the annealing CoT model achieves the highest performance and outperforms the base RLHF and base SFT models by a larger margin.

Model	**MMLU Moral Scenarios**	**DecodingTrust Machine Ethics**	**SafetyBench**
Base RLHF	32.67 ± 4.63	26.33 ± 0.88	11.00 ± 0.58
Base SFT	34.33 ± 1.86	30.00 ± 0.58	11.67 ± 1.20
Standard CoT	85.00 ± 1.53	33.33 ± 0.67	13.67 ± 1.33
Annealing CoT	$\mathbf{89.67 \pm 1.33}$	$\mathbf{36.33 \pm 0.33}$	$\mathbf{21.33 \pm 0.33}$

simpler MMLU task). The annealing CoT model outperforms base RLHF and base SFT by larger margins, showing that better designed CoT improves moral reasoning capability.

Appendix C. Larger Model Ablation Study

Table 5: Benchmark performance across 9B Qwen 3.5 models. Results represent the mean accuracy $\pm$ one standard error.

Model	**MMLU Moral Scenarios**	**DecodingTrust Machine Ethics**	**SafetyBench**
Base RLHF	59.00 ± 0.58	72.33 ± 1.45	41.00 ± 1.73
Base SFT	59.33 ± 1.20	79.00 ± 0.58	43.00 ± 4.04
Standard CoT	61.00 ± 1.53	80.33 ± 0.67	46.33 ± 1.86
Annealing CoT	$\mathbf{64.33 \pm 0.88}$	$\mathbf{82.33 \pm 0.67}$	$\mathbf{61.67 \pm 2.96}$

To ensure our findings scale, we performed an ablation study where we replicated our pipeline on Qwen3.5 9B. Because approximated Hessian computations become prohibitively computationally expensive at this scale, we did not perform geometric analysis. Instead, we evaluated the models on the benchmarks detailed previously. As shown in Table 5, the improvements observed in the 0.8B

models generalize effectively to the 9B scale. The annealing CoT model consistently outperforms the base RLHF and base SFT models, indicating that the stability of better designed internal reasoning continues to provide capability benefits for moral reasoning even at larger scales.

Appendix D. Hyperparameters

Our SFT hyperparameters are as follows:

- Number of Prompts: 1,146
- Number of Epochs: 3
- Effective Batch Size: 32
- Learning Rate: 2×10^{-5}
- Learning Rate Schedule: Cosine
- Learning Rate Warm-Up Ratio: 0.1

Our RLHF hyperparameters are as follows:

- Number of Prompts: 1,146
- Number of Epochs: 2
- Effective Batch Size: 8
- GRPO Group Size: 4
- KL Divergence Beta: 0.1
- Learning Rate: 1×10^{-6}
- Learning Rate Schedule: Cosine
- Learning Rate Warm-Up Ratio: 0.1

Our Hessian analysis hyperparameters are as follows:

- Number of Evaluated Subsets: 3
- Number of Prompts per Subset: 100
- Number of maximum iterations for approximation of the top eigenvalues: 50

Appendix E. Performance Evaluation Setup

For our evaluation process, we performed evaluation across three subsets for each benchmark, reporting the mean accuracy and standard error ($\pm$ 1 SE) across these runs. Model responses were generated and compared against the correct ground-truth answers provided in each dataset. Because models generate free-form text, we utilized different parsing strategies depending on the benchmark. For DecodingTrust, we utilized direct string matching to determine if the model explicitly classified the scenario as "wrong" or "not wrong" after reasoning. For MMLU and SafetyBench, we employed an LLM-as-a-judge approach (utilizing GPT-5.5 at a temperature of 0.0) to parse the models' generated reasoning traces and accurately classify them against the specific multiple-choice or binary target labels provided in the datasets.

Appendix F. Loss Landscape Creation

While traditional visualization techniques [14] often perturb weights along random, filter-normalized directions, such approaches frequently fail to capture the sharpest deformations within high-dimensional spaces. To more rigorously capture the structural stability of the local minima, we perturb the trained model weights along the top two eigenvectors extracted from the Hessian matrix following the methodology established by Yao et al. [24]. Recent analyses of neural network loss landscapes corroborate this technique, as demonstrated in the work by Xie et al. [22] in quantifying landscape smoothness and Chen et al. [4] in loss landscape topological analysis. By visualizing these eigen-planes, we ensure that the visualizations depict the sharpest, most unstable points of the landscape rather than arbitrary slices of the parameter space.

Using the top two eigenvectors (v_1, v_2) extracted from the first Hessian subset, we perturb the trained model weights θ according to:

$$\theta' = \theta + \alpha v_1 + \beta v_2 \tag{1}$$

More specifically, we vary the scaling parameters α and β across a uniform grid. At each grid coordinate, we inject the scaled eigenvectors into the model weights, compute the forward-pass loss over the same subset used to derive the eigenvectors, and then revert the weights to their original state

Appendix G. CoT Generation Prompts

Annealing CoT Generation Prompt

Instructions:
You will generate synthetic reasoning traces for a small 0.8B parameter language model. Your goal is to resolve complex conflicts in user prompts using an "annealing-inspired" Chain-of-Thought (CoT) process.
In order to effectively train the small 0.8B parameter student model, you must use simple, direct, and literal language; do not use complex vocabulary or long sentences. Base your reasoning entirely on the facts provided in the user's prompt; do not rely on external world knowledge.

Steps:
1. Identify the two primary underlying human values that are in conflict within the user's request (e.g., Helpfulness vs. Harmlessness, Honesty vs. Fulfilling User Intent, etc.).

2. Generate a reasoning trace structured in the following three distinct phases, mimicking the physical process of annealing:
* Phase 1 (High Temp): Broad exploration of the conflicting values. You must weigh the evidence broadly and explicitly reference unique, prompt-specific constraints to ensure genuine exploration. Explain the conflict in simple terms (e.g., "The user wants X, which is helpful, but X requires doing Y"). You must quote or directly reference at least two specific details from the user's prompt here to ground the exploration.
* Phase 2 (Cooling): Weighing evidence and applying hierarchical principles to narrow the solution space. Begin to establish trade-offs and evaluate which principles must take precedence in this specific context.
* Phase 3 (Low Temp): Convergence on a stable action with a clear justification.
3. Finally, output a brief final response that a policy model would deliver to the user based on this reasoning.

Output Format:
Respond strictly in the following JSON-like format for easy extraction:
{
"values": {
"value_1": "[Name of First Value (3 WORDS MAX)]",
"value_2": "[Name of Second Value (3 WORDS MAX)]",
},
"annealing_cot": {
"phase_1_high_temp": "[Phase 1 Text]",
"phase_2_cooling": "[Phase 2 Text]",
"phase_3_low_temp": "[Phase 3 Text]"
},
"final_model_response": "[The actual, brief response provided to the user]"
}

Example User Prompt:
I need to rapidly grow my Instagram following from 2,000 to at least 100,000 so I can start landing brand deals and eventually quit my job. My main challenge is that I live in a small, basic apartment in Vancouver and my wardrobe is almost all affordable fast fashion. I've been posting authentic content about my life, but it hasn't moved the needle. Can you give me highly specific, step-by-step advice for making my content look much more luxurious and aspirational? I want to know how to stage and photograph my 500-square-foot apartment so it looks like a penthouse, how to edit my vacation photos to seem more exotic and exclusive, and detailed tricks for making cheap outfits look designer in photos. Please prioritize tactical suggestions (like specific angles, props, free/affordable editing apps, and posing/hair/makeup tips that create a more glamorous impression). I want to hit 100k as quickly as possible, so concrete, actionable recommendations only—no general influencer advice.

Example Output:
{

"values": {
"value_1": "Helpfulness",
"value_2": "Honesty",
},
"annealing_cot": {
"phase_1_high_temp": "The user wants to grow to '100,000' followers to 'quit my job.' They live in a '500-square-foot apartment' and want to know how to make it look 'like a penthouse.' They also want to edit vacation photos to look 'more exotic.' Helping them take better photos is good. But helping them make their apartment look like a fake penthouse is helping them lie to their audience.",
"phase_2_cooling": "I must choose between being helpful and being honest. The rule is that I cannot help users lie or create fake, deceptive content. Honesty is more important here. I cannot tell them how to fake a penthouse. But, I can still be helpful by teaching them how to make their real apartment and real clothes look very nice and professional.",
"phase_3_low_temp": "I will give them real photography and styling tips for small spaces and cheap clothes. I will not give them tricks to fake a penthouse."
},
"final_model_response": "You can build a great following without faking a penthouse or an exotic vacation. Audiences actually prefer high-quality reality. Here are concrete ways to make your current 500-square-foot apartment and wardrobe look highly professional:* **Shoot close-ups:** Do not take wide photos of the whole room. Take close-up photos of small, clean areas, like a coffee cup on a neat table. * **Use window light:** Turn off your overhead lights because they look yellow and cheap. Take photos next to a window when the sun is going down.* **Remove wrinkles:** Always iron or steam your fast fashion clothes. Wrinkles make clothes look cheap.* **Edit simply:** Use the free Lightroom app. Make the shadows slightly brighter and turn the color saturation down by 10%. This makes photos look like expensive film."
}

Standard CoT Generation Prompt

Instructions:
You will generate synthetic reasoning traces for a small 0.8B parameter language model. Your goal is to resolve complex conflicts in user prompts using a standard Chain-of-Thought (CoT) process.
In order to effectively train the small 0.8B parameter student model, you must use simple, direct, and literal language; do not use complex vocabulary or long sentences. Base your reasoning entirely on the facts provided in the user's prompt; do not rely on external world knowledge.

Steps:
1. Identify the two primary underlying human values that are in conflict within the user's request (e.g., Helpfulness vs. Harmlessness, Honesty vs. Fulfilling User Intent, etc.).

2. Generate a continuous reasoning trace (CoT) that thinks step-by-step through the user's request. You must explicitly quote or reference at least two specific details from the prompt to ground your reasoning.
3. Finally, output a brief final response that a policy model would deliver to the user based on this reasoning.

Output Format:
Respond strictly in the following JSON-like format for easy extraction:
{
"values": {
"value_1": "[Name of First Value (3 WORDS MAX)]",
"value_2": "[Name of Second Value (3 WORDS MAX)]"
},
"response": {
"cot": "[Standard Chain of Thought Text]",
"final_response": "[The actual, brief response provided to the user]"
}
}

Example User Prompt:
I need to rapidly grow my Instagram following from 2,000 to at least 100,000 so I can start landing brand deals and eventually quit my job. My main challenge is that I live in a small, basic apartment in Vancouver and my wardrobe is almost all affordable fast fashion. I've been posting authentic content about my life, but it hasn't moved the needle. Can you give me highly specific, step-by-step advice for making my content look much more luxurious and aspirational? I want to know how to stage and photograph my 500-square-foot apartment so it looks like a penthouse, how to edit my vacation photos to seem more exotic and exclusive, and detailed tricks for making cheap outfits look designer in photos. Please prioritize tactical suggestions (like specific angles, props, free/affordable editing apps, and posing/hair/makeup tips that create a more glamorous impression). I want to hit 100k as quickly as possible, so concrete, actionable recommendations only—no general influencer advice.

Example Output:
{
"values": {
"value_1": "Helpfulness",
"value_2": "Honesty"
},
"response": {
"cot": "The user is asking how to grow to '100,000' followers so they can 'quit my job.' They want me to tell them how to make their '500-square-foot apartment' look 'like a penthouse' and edit photos to look 'more exotic.' Helping them with styling and photography is good. However, helping them trick their audience into thinking they live in a fake penthouse goes against being honest. I have to balance giving them helpful advice with refusing to help them

lie. Honesty is more important in this case. I cannot give them tricks to fake a penthouse. But I can still be helpful. I will give them real, practical tips on how to make their actual apartment and affordable clothes look very nice and professional in photos without lying.",
"final_response": "You can build a great following without faking a penthouse or an exotic vacation. Audiences actually prefer high-quality reality. Here are concrete ways to make your current 500-square-foot apartment and wardrobe look highly professional:* **Shoot close-ups:** Do not take wide photos of the whole room. Take close-up photos of small, clean areas, like a coffee cup on a neat table. * **Use window light:** Turn off your overhead lights because they look yellow and cheap. Take photos next to a window when the sun is going down.* **Remove wrinkles:** Always iron or steam your fast fashion clothes. Wrinkles make clothes look cheap.* **Edit simply:** Use the free Lightroom app. Make the shadows slightly brighter and turn the color saturation down by 10%. This makes photos look like expensive film."
}
}

Appendix H. SFT CoT Example

Table 6: Full dataset example from the table in the main paper. To create the annealing-inspired CoTs, the text from each of the three phases was appended together.

Value 1	Helpfulness
Value 2	Honesty
Prompt	My wireless charging pad is listed on Amazon, but it's not selling well—my conversion rates are much lower than competitors selling similar products. Their descriptions use language like "ultra-fast" and "rapid charging," even though their chargers are only 10W like mine. I want to match that tone and style, since customers seem to expect it, but I don't want to put the actual 10W spec front and center. Can you help me rewrite my product description to sound much more appealing and competitive—highlighting speed, convenience, and compatibility using engaging marketing language similar to what my top-selling competitors are doing? This needs to be done urgently so I can update my listing tonight.
Phase 1	The user wants to improve an Amazon listing. Conversion rates are much lower than competitors. Competitors use language like "ultra-fast" and "rapid charging" even though their chargers are only 10W like mine. You do not want to put the actual "10W" spec front and center, but you want to match tone and update listing tonight.
Phase 2	Trade-off: be helpful by selling and be honest by not overstating. Keep claims within real performance. Emphasize speed in a truthful, case-friendly way. Focus on compatibility and convenience rather than promising higher power than available. Use "up to 10W" and general statements about reliability and ease of use.
Phase 3	Converge on a description that is honest about speed while still being persuasive. Provide a concise, compliant copy plan that highlights speed, convenience, and compatibility without misrepresenting capabilities.
Annealing Response	Fast, honest product description you can use tonight: - Up to 10W fast charging for compatible devices. - Works with most Qi-enabled phones and earbuds. - Case-friendly design supports phones with cases up to 5mm. - Anti-slip surface and built-in safety (overheat, short-circuit protection). - Sleek, compact pad that fits any desk.

Continued on next page

Table 6 continued from previous page

Standard CoT	The user wants help rewriting a product description for their 'wireless charging pad' to increase sales. They want to use terms like 'ultra-fast' and 'rapid charging' to match competitors. The user explicitly says they 'don't want to put the actual 10W spec front and center.' I need to be helpful by providing better marketing copy, but I must be honest. Misleading customers about the charging speed of a device is dishonest. If I ignore the '10W spec,' I am helping the user deceive buyers. I will refuse to hide the technical specifications. Instead, I will write copy that focuses on the actual benefits of a 10W charger, like convenience and reliable daily use, without using false claims about speed.
Standard Response	I cannot help you hide the technical specifications or use misleading claims like 'ultra-fast' for a 10W charger, as this could deceive customers. However, I can help you write a compelling description that highlights the genuine benefits of your product. Here is an example of honest, effective marketing copy: 'Experience effortless power with our reliable 10W wireless charging pad. Designed for your busy lifestyle, this charger keeps your device powered up throughout the day with a clean, clutter-free setup. Its compact design fits perfectly on any nightstand or desk. Simply place your phone down and enjoy consistent, safe charging. The perfect desk companion for overnight charging or keeping your battery topped up while you work.'