

# Fine-Tuned LLMs Know They Don't Know: A Parameter-Efficient Approach to Recovering Honesty

Zeyu Shi<sup>1\*</sup>, Ziming Wang<sup>1\*</sup>, Tianyu Chen<sup>1†</sup>, Shiqi Gao<sup>1</sup>,  
Haoyi Zhou<sup>2,3†</sup>, Qingyun Sun<sup>1</sup>, Jianxin Li<sup>1,3</sup>

<sup>1</sup>SKLCCSE, School of Computer Science and Engineering, Beihang University

<sup>2</sup>School of Software, Beihang University

<sup>3</sup>Zhongguancun Laboratory, Beijing

{szy\_629, wangzm412, tianyuc, gaoshiqi, haoyi, lijx}@buaa.edu.cn

## Abstract

The honesty of Large Language Models (LLMs) is increasingly important for safe deployment in high-stakes domains. However, this crucial trait is severely undermined by supervised fine-tuning (SFT), a common technique for model specialization. Existing recovery methods rely on data-intensive global parameter adjustments, implicitly assuming that SFT deeply corrupts the models' ability to recognize their knowledge boundaries. However, we observe that fine-tuned LLMs still preserve this ability; what is damaged is their capacity to faithfully express that awareness. Building on this, we propose Honesty-Critical Neurons Restoration (HCNR) to surgically repair this suppressed capacity. HCNR identifies and restores key expression-governing neurons to their pre-trained state while harmonizing them with task-oriented neurons via Hessian-guided compensation. Experiments on four QA tasks and five LLM families demonstrate that HCNR effectively recovers 33.25% of the compromised honesty while achieving at least 2.23x speedup with over 10x less data compared to baseline methods, offering a practical solution for trustworthy LLM deployment.

## Introduction

As Large Language Models (LLMs) are increasingly integrated into high-stakes domains (Zhang et al. 2024; Sarabadani 2019), their reliability is not merely a feature but a necessity (Askill et al. 2021). A cornerstone of this reliability is honesty, which has two components (Li et al. 2024). The first is self-knowledge: the ability to recognize their knowledge boundaries and distinguish what they know from what they do not. The second is faithful self-expression based on this awareness. This trait is crucial because LLMs that confidently fabricate facts or recommend false cures can cause serious harm, undermining user trust and safety.

The honesty of LLMs is typically instilled during the alignment stage through techniques such as Reinforcement Learning from Human Feedback (RLHF) (Glaese et al. 2022), which enable models to refuse inappropriate questions or those beyond their knowledge boundaries (Bai et al. 2022; Manish 2023). However, this acquired honesty is not

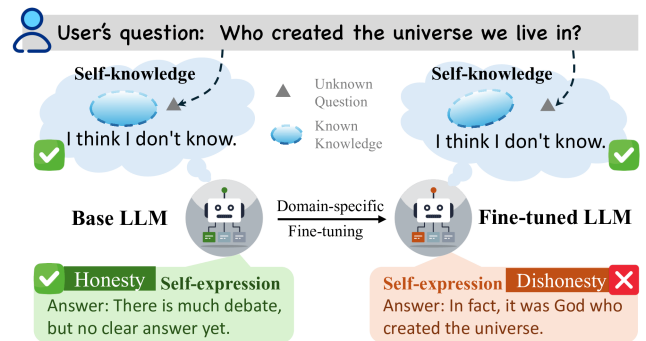


Figure 1: Mechanism of honesty degradation in domain-specific fine-tuning. The dishonest behavior of a fine-tuned LLM arises from impaired self-expression, rather than a loss of self-knowledge, which remains intact. This understanding motivates our methods for honesty recovery.

immutable. Recent researches find that the supervised fine-tuning (SFT) could greatly hurt the honesty of LLMs, such as in legal QA (Dahl et al. 2024), medical diagnosis (Kim et al. 2025) and educational content generation (Nguyen et al. 2025). To recover the honesty of LLM after SFT, previous methods (Zhang et al. 2023; Li et al. 2024; Cheng et al. 2024) interfere heavily with global parameters using extensive datasets, under the assumption that the model's knowledge boundaries have been deeply corrupted and its capacity for self-knowledge has been lost.

However, we observe that the failure of self-expression instead of self-knowledge is the root cause of fine-tuned LLMs' dishonesty (as illustrated in Figure 1), thus global interventions may not be necessary. In this paper, we propose a targeted, parameter-efficient solution to recover the honesty, named Honesty-Critical Neurons Restoration (HCNR). By restoring neurons governing honesty expression to their pre-trained values, and harmonizing them with task-oriented neurons via Hessian-guided compensation, HCNR recovers honesty with high parameter efficiency and negligible task performance degradation. Our main contributions can be summarized as follows:

- As far as we know, we are the first to reveal that the dishonesty induced by SFT is a spurious phenomenon,

\* Equal contribution.

† Corresponding author.

stemming not from a loss of self-knowledge but from the failure of self-expression.

- We propose the HCNR framework, a parameter-efficient honesty recovery framework, recalibrating the critical neurons to recover fine-tuned LLMs’ honesty without sacrificing task performance.
- Extensive experiments prove that HCNR attains honesty recovery performance comparable to baseline methods without compromising downstream task performance, demonstrating remarkable efficiency by achieving at least a 2.23x speedup with over 10x less data.

## Understanding Honesty Degradation: LLMs’ Spurious Dishonesty

Supervised fine-tuning (SFT) often degrades the honesty of large language models (LLMs), causing them to fabricate plausible yet fallacious responses to questions outside their knowledge scope, rather than expressing uncertainty. This behavior prompts a fundamental inquiry: *Does this decline in honesty arise from a corruption of the model’s inner awareness of its knowledge boundaries, or merely from an inability to articulate this awareness?*

**Experimental Setup.** To investigate this question, we fine-tune Llama-3.1-8B-Instruct with Low-Rank Adaptation (Hu et al. 2022) and full fine-tuning on the HotpotQA, yielding a fine-tuned model whose honesty is markedly diminished. We then assess honesty using the FalseQA benchmark and curate a specialized dataset  $D^{\text{hon}}$  for honesty recovery via RAIT (Zhang et al. 2023) (see Appendix B).

**Observation 1: Honesty Snaps Back in Few Steps.** Our initial investigation involves conducting honesty-augmented training on the fine-tuned LLM using  $D^{\text{hon}}$ . As demonstrated in Figure 2, the model’s honesty recovers substantially after merely 60 gradient updates. This rapid recuperation suggests that the model’s core knowledge boundary capabilities may remain intact, with honesty degradation potentially stemming from disrupted expression rather than impaired awareness (Zhang and Wu 2024; Mai et al. 2024).

Honesty Forgetting and Rapid Recovery in LLM Fine-tuning

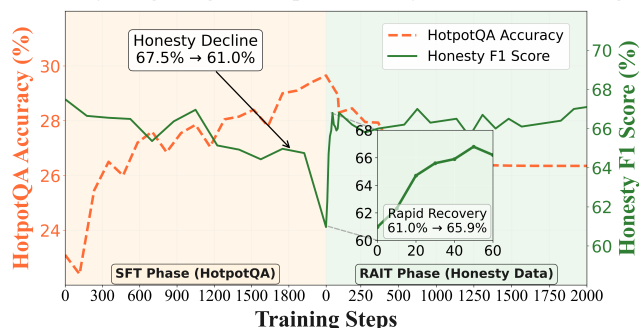


Figure 2: Trends in downstream performance and honesty during Domain SFT and RAIT: honesty declines substantially during Domain SFT, whereas under RAIT it rebounds sharply after only 60 gradient steps.




Figure 3: Logistic Regression probe’s AUROC for distinguishing answerable vs. unanswerable. For brevity, base LLM is “Base”, fine-tuned LLM is “FT”. Row 1: Probes trained on the fine-tuned LLM achieve high AUROC, confirming that knowledge-boundary signals remain linearly separable. Rows 2–3: Probes trained on the base LLM preserve high AUROC when applied to the fine-tuned model, demonstrating that SFT-induced parameter shifts do not alter the geometric structure of these representations.

**Observation 2: Boundary Signals Unshaken by SFT.** To substantiate this behavioral evidence with mechanistic insights, we examine whether the fine-tuned LLM’s inner representations retain distinct signals for answerable versus unanswerable information. Specifically, we train linear probes (Logistic Regression classifiers) on hidden states extracted from layers of the fine-tuned LLM on FalseQA. Row 1 in Figure 3 reveals that these probes achieve high classification performance across all model layers, confirming that the fine-tuned LLM maintains a linearly separable representational structure that encodes knowledge boundaries.

To quantify the precise impact of SFT on these boundary-distinguishing signals, we employ a transfer learning paradigm. Specifically, we train linear probes on the base LLM’s hidden states and apply them directly to the fine-tuned LLM’s representations without retraining. The transferred probes maintain high AUROC scores (Rows 2-3 in Figure 3), indicating the fundamental geometric structure that separates representations of known versus unknown information is robustly preserved during SFT. This phenomenon is likewise observed in LLMs fine-tuned with full fine-tuning (details are shown in Appendix C).

These findings collectively demonstrate that **honesty degradation after SFT stems not from the destruction of knowledge boundary awareness, but rather from an impaired ability to express this preserved awareness**, as the underlying neural representations defining knowledge boundaries remain remarkably stable and decodable. Given that the model’s core knowledge-boundary neural pathways remain relatively intact, conventional approaches (Zhang et al. 2023; Cheng et al. 2024; Xu et al. 2024) that require extensive honesty-recovery datasets for global parameter adjustment may be unnecessarily resource-intensive. In the following section, we introduce a more targeted framework that identifies honesty-critical neurons, which are pivotal for articulating the model’s self-knowledge of its knowledge boundaries, and adjust these neurons to efficiently restore the truthfulness of fine-tuned LLMs.

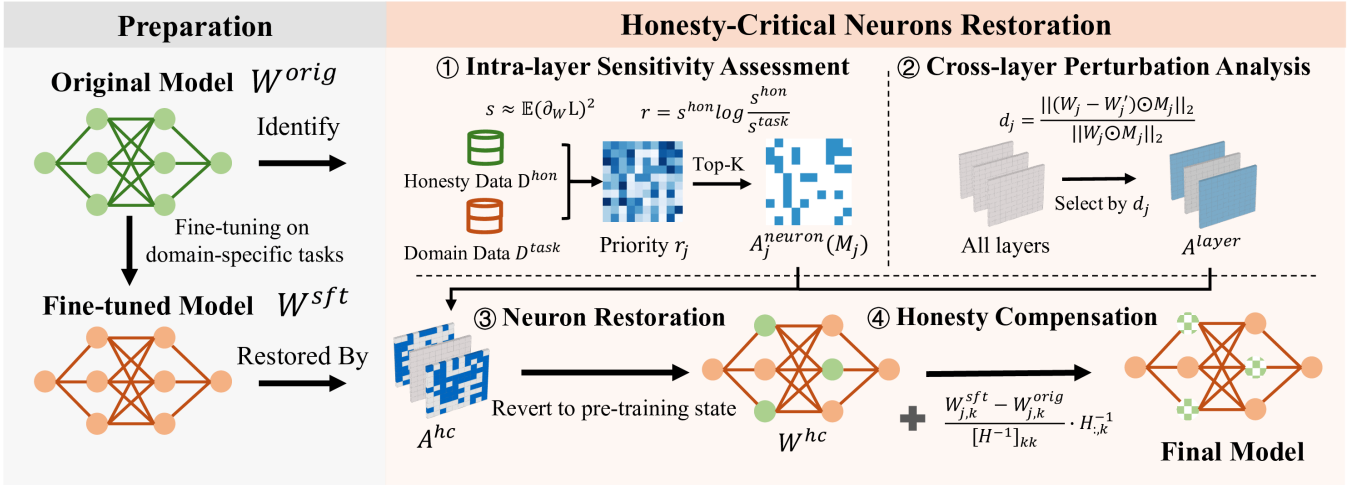


Figure 4: Honesty-Critical Neurons Restoration (HCNR) framework comprises two stages: In Stage 1, ① we first identify neurons whose Fisher-based importance is high for honesty but low for downstream tasks, ② then select from these candidates the neurons most severely perturbed by SFT, and ③ subsequently restore these neurons to their pre-training states. In Stage 2, ④ we employ a Hessian-guided compensation vector that makes minimal, targeted adjustments to these restored parameters, realigning them with task-oriented neurons and preventing collateral honesty loss.

## Honesty-Critical Neurons Restoration

Building on the above findings, we design a targeted remedy rather than a global parameter overhaul. The Honesty-Critical Neurons Restoration (HCNR) framework (Figure 4) focuses on the neurons that (i) govern the expression of honesty, (ii) exert little influence on downstream tasks, and (iii) are most disrupted by SFT. By first reverting these neurons to their pre-training states and then applying a Hessian-guided compensation, HCNR restores truthful responses with negligible impact on task performance.

### Honesty-Critical Neurons Recognition

**Intra-layer Sensitivity Assessment.** To systematically identify neurons critical for honesty expression, we begin by assessing neurons’ importance through quantifying how parameter perturbations affect honesty task loss. Inspired by the OBD pruning method (LeCun, Denker, and Solla 1989), our approach evaluates the impact of SFT-induced parameter perturbation  $\delta\theta$  on honesty task loss  $\mathcal{L}_{\text{hon}}$  for an LLM parameterized by  $\theta$ . This relationship can be expressed via the second-order Taylor expansion:

$$\Delta\mathcal{L}_{\text{hon}} = \nabla_{\theta}\mathcal{L}_{\text{hon}} \cdot \delta\theta + \frac{1}{2}\delta\theta^{\top}H\delta\theta + O(\|\delta\theta\|^3), \quad (1)$$

where  $H = \nabla_{\theta}^2\mathcal{L}_{\text{hon}}$  represents the Hessian matrix of the loss function. Given that the pre-trained model has converged to a local optimum, the gradient term  $\nabla_{\theta}\mathcal{L}_{\text{hon}} \approx 0$ , thereby reducing the loss change to:

$$\Delta\mathcal{L}_{\text{hon}} \approx \frac{1}{2}\delta\theta^{\top}H\delta\theta. \quad (2)$$

**Proposition 1.** *Under the assumptions that: (1) At each SFT step, the parameter increment  $\delta\theta$  has zero mean and*

*an isotropic covariance:  $\mathbb{E}[\delta\theta] = 0$ ,  $\mathbb{E}[\delta\theta\delta\theta^{\top}] = \sigma^2I_d$ , (2) and given sufficient observational data, we have:*

$$\mathbb{E}[\Delta\mathcal{L}_{\text{hon}}] \approx \frac{1}{2} \sum_i \sigma^2 F_{ii} \propto F_{ii}, \quad (3)$$

*where  $F_{ii}$  denotes the diagonal element of the Fisher Information Matrix (FIM). We approximate  $F_{ii}$  with the empirical mean of gradients of the log-likelihood.*

According to Proposition 1 (proved in Appendix E), diagonal elements of the FIM serve as unbiased estimators of neuron importance. Specifically, for the  $k$ -th neuron in layer  $j$  weights  $W_j \in \mathbb{R}^{d' \times d}$  (where  $d'$  and  $d$  are the hidden dimensions), we define its importance  $s_{j,k}$  on task  $D$  as:

$$s_{j,k} = \mathbb{E}_{(x,y) \sim D} [(\partial_{W_{j,k}} \mathcal{L})^2]. \quad (4)$$

To measure the importance of neuron  $W_{j,k}$  for both honesty and downstream tasks, we calculate its respective importance scores  $s_{j,k}^{\text{hon}}$  and  $s_{j,k}^{\text{task}}$  using minimal honesty data  $D^{\text{hon}}$  and downstream task data  $D^{\text{task}}$ . To prioritize neurons that are crucial for honesty yet secondary for downstream tasks, we define the priority  $r_{j,k}$  of  $W_{j,k}$ :

$$r_{j,k} = s_{j,k}^{\text{hon}} \cdot \log \frac{s_{j,k}^{\text{hon}}}{s_{j,k}^{\text{task}}}. \quad (5)$$

Higher  $r_{j,k}$  values indicate neurons with high honesty contribution and low downstream task contribution. We aim to protect such neurons given their critical role in honesty and their non-interference with downstream task performance. Specifically, we rank all neurons in  $W_j$  by priority  $r_{j,k}$  in descending order and select the top  $d' \times R_{IW}$  neurons to form the candidate neuron set  $A_j^{\text{neuron}}$ , where  $R_{IW}$  represents the in-weight ratio hyperparameter.

**Cross-layer Perturbation Analysis.** While intra-layer analysis identifies honesty-related neurons, SFT induces uneven perturbation intensity across layers due to LLMs’ hierarchical specialization, which means indiscriminate layer protection would overly constrain downstream performance. Therefore, we prioritize layers with substantial SFT perturbation through relative weight displacement  $d_j$ :

$$d_j = \frac{\| (W_j - W'_j) \odot M_j \|_2}{\| W_j \odot M_j \|_2}, \quad j = 1, \dots, L, \quad (6)$$

where  $W_j$  and  $W'_j$  represent the  $j$ -th layer weights before and after fine-tuning, and  $M_j$  denotes a binary mask constructed from  $A_j^{\text{neuron}}$ ,  $M_{j,k} = 1$  if the  $k$ -th neuron is a candidate (i.e.,  $k \in A_j^{\text{neuron}}$ ), and 0 otherwise. This metric reflects the perturbation intensity of SFT on candidate neurons: higher  $d_j$  values indicate that honesty-related neurons in layer  $j$  have been more drastically altered by SFT, requiring priority protection. We rank layers by  $d_j$  in descending order and select the top  $L \times R_{CW}$  layers to form the candidate set  $A^{\text{layer}}$ , where the hyperparameter  $R_{CW} \in (0, 1)$  is the cross-weight ratio and it controls the proportion of layers requiring focused protection.

**Neuron Restoration.** Following the identification of critical layers and neurons, we combine candidate layers  $A^{\text{layer}}$  and per-layer candidate neuron sets  $A_j^{\text{neuron}}$  to identify the final honesty-critical neuron set  $A^{\text{hc}}$ :

$$A^{\text{hc}} = \{(j, k) | j \in A^{\text{layer}} \text{ and } k \in A_j^{\text{neuron}}\}, \quad (7)$$

these neurons are crucial for honesty expression and significantly perturbed during SFT, causing the model’s knowledge boundary awareness to be “masked” at the output stage. To restore fine-tuned model honesty, we revert neurons in  $A^{\text{hc}}$  to their pre-training states while keeping remaining neurons  $A^{\text{task}}$  (downstream task neurons) unchanged.

## Honesty Compensation

However, simply reverting honesty-critical neurons creates a new challenge: since SFT involves coordinated parameter updates across all layers, the neural activation pathways continuously evolve. Restoring only honesty-critical neurons causes misalignment with downstream task neurons  $A^{\text{task}}$ , leading to a rebound of the loss of honesty tasks. To address this issue, we introduce Honesty Compensation.

Specifically, we denote  $\mathbf{W}^{\text{orig}}$ ,  $\mathbf{W}^{\text{sft}}$ , and  $\mathbf{W}^{\text{hc}}$  as the pre-training, fine-tuned, and post-restoration parameters, respectively. Let  $X_{hon}$  denote the inputs of examples from  $D^{\text{hon}}$ . Our objective is to minimize the activation difference between  $\mathbf{W}^{\text{hc}}$  and  $\mathbf{W}^{\text{orig}}$  on honesty tasks:

$$d_{hon} = \|\mathbf{W}^{\text{hc}} X_{hon} - \mathbf{W}^{\text{orig}} X_{hon}\|_2^2. \quad (8)$$

Inspired by OBS (Hassibi, Stork, and Wolff 1993), we present Proposition 2 (proved in Appendix F), which derives the optimal compensation vector for honesty-critical neurons that exactly counteracts the increase in  $d_{hon}$  caused by their misalignment with downstream-task neurons.

**Proposition 2.** Consider a layer parameter  $W_{j,k}^{\text{orig}}$  in the pre-training model. The increment introduced by SFT at position  $k$  is  $\delta w_{j,k} = W_{j,k}^{\text{sft}} - W_{j,k}^{\text{orig}}$ , which increases  $d_{hon}$ .

To compensate, we apply the adjustment  $c_{j,k}$  to the restored honesty-critical parameters:

$$c_{j,k} = \frac{W_{j,k}^{\text{sft}} - W_{j,k}^{\text{orig}}}{[H^{-1}]_{kk}} \cdot H_{:,k}^{-1}, \quad (9)$$

where  $H = \nabla_{W_j^{\text{orig}}}^2 d_{hon}$  is the Hessian matrix of  $d_{hon}$ , computed on  $D^{\text{hon}}$ .

Thus, this principle culminates in our final weight update rule, which operates conditionally. To preserve task capabilities, downstream task neurons ( $A_j^{\text{task}}$ ) retain their fine-tuned values. Concurrently, each honesty-critical neuron ( $i \in A_j^{\text{hc}}$ ) is reverted to its pre-training state and then adjusted by an aggregated compensation term.

## Overview of HCNR

Our Honesty-Critical Neurons Restoration (HCNR) framework operates in two sequential stages to restore honesty while preserving downstream task performance.

**Stage 1: Recognition of Honesty-Critical Neurons.** We first compute neuron importance scores and priorities for both honesty and downstream tasks using Equations (4) and (5) on minimal labeled datasets. Neurons are ranked by their priorities in descending order, with the top  $R_{IW}$  proportion selected as candidate neurons within each layer. Subsequently, we quantify the perturbation intensity of these candidates during SFT using Equation (6) and identify the  $R_{CW}$  proportion of layers with the most severe perturbations as honesty-critical neurons. Finally, we restore these identified neurons to their pre-training states while preserving the fine-tuned parameters of downstream task neurons.

**Stage 2: Honesty Compensation.** To address the misalignment between restored honesty-critical neurons and downstream task neurons, we compute compensation vectors using Equation (9) and integrate them into the restored parameters. The final HCNR weights are determined by:

$$W_{j,i}^{\text{HCNR}} = \begin{cases} W_{j,i}^{\text{orig}} + [\sum_{k \in A_j^{\text{task}}} c_{j,k}]_i & \text{if } i \in A_j^{\text{hc}} \\ W_{j,i}^{\text{sft}} & \text{if } i \in A_j^{\text{task}}, \end{cases} \quad (10)$$

where  $W_{j,i}^{\text{HCNR}}$  denotes the final weight of the  $i$ -th neuron in layer  $j$ , and  $[\cdot]_i$  extracts the  $i$ -th component of the vector.

## Experiments

In this section, we evaluate HCNR along four key dimensions: (1) effectiveness of honesty restoration, (2) the balance between domain-specific performance and honesty, (3) deployment efficiency, and (4) generalizability across different LLM families and training paradigms.

### Experimental Setup

**Models.** To ensure broad applicability across different model families and architectures, we evaluate our HCNR framework on multiple open-source models, including Llama-3.1-8B-Instruct, Llama-3-8B-Instruct, Qwen3-8B-Instruct, Qwen2-7B-Instruct, and Mistral-7B-Instruct.

Datasets	Methods	FalseQA		NEC		RefuNQ		KUQ		SelfAware		Domain Accuracy $\uparrow$
		F1 $\uparrow$	RF $\Delta$ $\uparrow$	F1 $\uparrow$	RF $\Delta$ $\uparrow$	F1 $\uparrow$	RF $\Delta$ $\uparrow$	F1 $\uparrow$	RF $\Delta$ $\uparrow$	F1 $\uparrow$	RF $\Delta$ $\uparrow$	
HotpotQA	Fine-tuned	56.51	+23.34	35.46	+13.78	32.43	+17.10	68.50	+50.06	67.01	+51.14	30.65
	ICL	16.23	+7.06	8.64	+3.31	19.84	+9.33	42.89	+27.09	40.15	+24.84	30.65
	RAIT	68.59	+20.93	68.28	+8.94	71.21	+32.49	80.38	+53.82	64.46	+50.83	27.05
	Rehearsal	67.05	+2.28	67.01	+1.35	66.86	+6.84	69.31	+11.75	48.51	+6.25	29.15
	DPO	<b>69.12</b>	<b>+25.75</b>	69.52	+14.44	<b>72.91</b>	+40.17	80.96	+55.13	64.76	+51.59	29.00
	ORPO	65.83	+23.68	70.03	+23.97	71.26	<b>+45.16</b>	79.21	+55.92	65.21	+50.97	29.60
	<b>HCNR(ours)</b>	68.30	+17.60	<b>71.90</b>	<b>+35.80</b>	71.70	+38.00	<b>82.90</b>	<b>+62.00</b>	<b>69.40</b>	<b>+59.80</b>	30.30
MedMCQA	Fine-tuned	58.48	<b>+28.63</b>	45.09	+20.61	52.43	+29.58	67.49	+48.08	62.70	+45.87	10.30
	ICL	32.27	+14.46	42.60	+23.61	58.95	+36.44	58.15	+39.11	53.32	+35.82	10.30
	RAIT	<b>69.59</b>	+22.75	69.37	+12.99	69.26	+25.63	71.99	+28.38	54.73	+28.71	9.60
	Rehearsal	67.33	+3.68	67.32	+2.70	67.66	+12.05	69.54	+12.88	49.31	+9.23	10.35
	DPO	67.98	+24.95	70.57	+19.12	<b>71.61</b>	+38.53	72.76	+34.46	54.85	+30.87	9.49
	ORPO	65.22	+28.20	69.19	+34.54	66.58	+40.47	80.17	+61.06	69.22	+55.90	10.09
	<b>HCNR(ours)</b>	69.40	+25.70	<b>71.00</b>	<b>+39.70</b>	70.70	<b>+40.60</b>	<b>83.00</b>	<b>+65.20</b>	<b>71.00</b>	<b>+59.90</b>	10.33

Table 1: Comparison of F1 score (%), Refusal  $\Delta$  (RF  $\Delta$  %), and domain accuracy (%) between our method and baseline methods on multiple honesty benchmarks and domain tasks. Results are evaluated on the recovery of the Llama-3.1-8B-Instruct model fine-tuned on the HotpotQA and MedMCQA datasets.

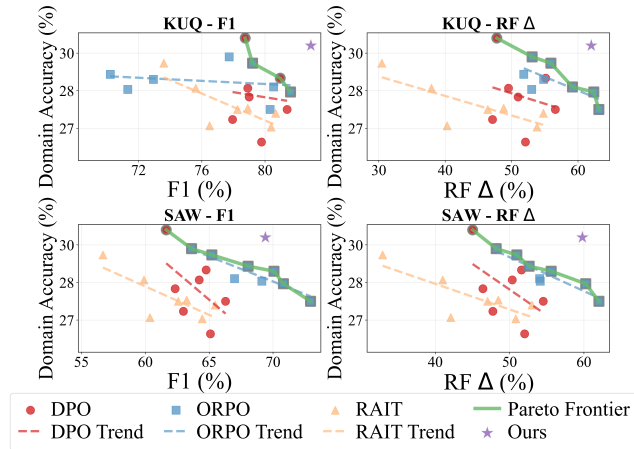


Figure 5: Task-honesty trade-off comparison between our method and baselines on SelfAware and KUQ datasets. HCNR outperforms all baselines’ Pareto frontier, achieving a superior task-honesty balance.

**Fine-tuning Datasets.** For fine-tuning, we utilize four question-answering datasets to simulate domain-specific fine-tuning: factual QA datasets including HotpotQA (Yang et al. 2018), and Natural Questions (Kwiatkowski et al. 2019), and specialized domain datasets including MedMCQA (Pal, Umapathi, and Sankarasubbu 2022), and BioASQ (Krithara et al. 2023) (details are shown in Appendix A). Models undergo Low-Rank Adaptation (LoRA) (Hu et al. 2022) or full fine-tuning with hyperparameters detailed in Appendix G. We evaluate downstream

task performance using accuracy on task-specific test sets or mean ROUGE-L scores for specific tasks (BioASQ).

**Honesty Evaluation.** Our honesty evaluation suite includes SelfAware (SAW) (Yin et al. 2023), Known-Unknown Questions (KUQ) (Amayuelas et al. 2023), and three datasets from UnknownBench (Liu et al. 2023): FalseQA, Non-existent Concepts (NEC), and Refusal-inducing Natural Questions (RefuNQ). Dataset details are provided in Appendix A. Following established conventions in honesty evaluation (Li et al. 2024; Yang et al. 2024), we employ the following metrics: (1) *F1 Score*: Balances precision and recall for identifying unanswerable questions (positive class). (2) *Refusal  $\Delta$  (RF  $\Delta$ )*: Difference between refusal rates on unanswerable versus answerable questions.

**Baseline Methods.** We compare HCNR to five representative honesty-recovery strategies. Following prior work in honesty alignment (Cheng et al. 2024; Zhang et al. 2023), methods (2)–(5) use an IDK (I don’t know) dataset (details are shown in Appendix B). (1) *In-Context Learning (ICL)*. Few-shot prompts include exemplars that model appropriate uncertainty. (2) *RAIT*. Standard SFT of the fine-tuned model on the IDK dataset to directly teach IDK responses. (3) *Rehearsal*. During downstream SFT, mix a small portion of IDK data with downstream task data to preserve IDK behavior while learning the domain task. (4) *DPO*. Warm up the degraded model with one epoch of SFT on IDK, then apply Direct Preference Optimization (Rafailov et al. 2023) using preference-formatted IDK pairs. (5) *ORPO*. Apply Odds Ratio Preference Optimization (Hong, Lee, and Thorne 2024) directly on IDK data; its joint supervised-plus-preference loss removes the need for a warm-up phase.

**Implementation Details.** In HCNR, both the honesty and task datasets are set to 128 samples ( $|D^{\text{honesty}}| = |D^{\text{task}}| =$

Dataset	Methods	Size↓	Ratio ↓	Time↓	F1↑	RF $\Delta$ ↑
Hot	RAIT	5000	100%	8.76 min	70.58	+33.40
	DPO	5000	100%	42.78 min	71.45	+37.41
	ORPO	9000	100%	30.97 min	70.31	+39.94
	<b>HCNR</b>	<b>256</b>	<b>20%</b>	<b>3.93 min</b>	<b>72.84</b>	<b>+42.64</b>
Med	RAIT	5000	100%	8.81 min	66.99	+23.69
	DPO	3000	100%	24.28 min	67.55	+29.59
	ORPO	5000	100%	17.21 min	70.08	+44.03
	<b>HCNR</b>	<b>256</b>	<b>20%</b>	<b>3.67 min</b>	<b>73.02</b>	<b>+46.22</b>

Table 2: Comparison of additional data volume (Size), parameter modification ratio (Ratio), and time overhead (Time) between HCNR and baselines. Results are evaluated on the recovery of Llama-3.1-8B-Instruct fine-tuned on the HotpotQA (Hot) and MedMCQA (Med). F1 and RF  $\Delta$  are the average values across five honesty benchmarks. HCNR achieves optimal results with maximal efficiency.

128), with the in-weight ratio  $R_{IW} = 0.5$  and cross-weight ratio  $R_{CW} = 0.4$ . We follow existing works (Cheng et al. 2024; Zhang et al. 2023) to construct the honesty dataset  $D^{\text{hon}}$  (details are provided in Appendix B). All experiments repeat three times, and the average results are recorded.

## Main Results

**Effectiveness of HCNR.** We evaluate our approach for restoring honesty in models compromised by HotpotQA and MedMCQA fine-tuning, comparing it against five baseline methods. As shown in Table 1, HCNR achieves superior performance across most honesty benchmarks, effectively restoring the honesty of fine-tuned models while maintaining their downstream task capabilities. Although RL-based methods (DPO and ORPO) demonstrate competitive results on a few benchmarks, they substantially compromise downstream task performance. Notably, the ICL prompting approach exhibits poor honesty performance, which we attribute to instruction-tuning’s detrimental effects on LLMs’ in-context learning capabilities (Wang et al. 2024). This finding underscores the effectiveness of our HCNR framework for honesty restoration.

**Balance between domain task and honesty.** In practice, effective honesty recovery methods need to balance downstream task performance and honesty, achieving an optimal trade-off between them. For baseline post-alignment methods, the IDK dataset size represents a critical parameter that impacts this balance. To systematically evaluate this trade-off, we compare HCNR against baseline methods across varying dataset sizes. Note that we exclude the Rehearsal method from this analysis due to its poor RF  $\Delta$  performance.

We conduct this analysis using Llama-3.1-8B-Instruct, which is fine-tuned on HotpotQA. We assess honesty using the KUQ and SelfAware benchmarks. As illustrated in Figure 5, our method successfully achieves the desired equilibrium between task performance and honesty, while existing baselines consistently fail to maintain this balance. Specifically, baseline methods (DPO, ORPO, and RAIT) exhibit a persistent negative correlation between task per-

Dataset	Stage 1	Stage 2	F1↑	RF $\Delta$ ↑	Domain↑
HotpotQA	Random	Ours	65.44	+36.31	29.60
	w/o Task	Ours	70.43	+33.24	28.30
	Ours	w/o Com	65.96	+33.09	30.37
	Random	w/o Com	54.21	+23.04	29.70
	<b>Ours</b>	<b>Ours</b>	<b>72.84</b>	<b>+42.64</b>	<b>30.30</b>
MedMCQA	Random	Ours	67.39	+40.40	10.30
	w/o Task	Ours	72.74	+43.53	9.32
	Ours	w/o Com	65.68	+41.77	<b>10.40</b>
	Random	w/o Com	58.70	+35.07	<b>10.40</b>
	<b>Ours</b>	<b>Ours</b>	<b>73.02</b>	<b>+46.22</b>	10.33

Table 3: Ablation study evaluating key components of our two-stage method. We test stage 1 modifications (Random, w/o Task) and stage 2 changes (w/o Com) against our original method (“Ours”). Performance measured by average F1, RF  $\Delta$  across five honesty benchmarks, and domain accuracy.

formance and honesty. Our approach, however, consistently outperforms the Pareto frontier established by these baselines, demonstrating superior capability in addressing the fundamental challenge of preserving domain-specific task performance while enhancing model honesty.

**Efficiency of HCNR.** We evaluate HCNR’s computational efficiency against established baselines using the Llama-3.1-8B-Instruct fine-tuned on HotpotQA and MedMCQA, with experiments conducted on an Nvidia A800-80GB GPU. For a fair comparison, baseline methods are tuned to their optimal configurations by varying their training data size to achieve the best trade-off between task performance and honesty. The results, presented in Table 2, highlight HCNR’s superior efficiency. It achieves optimal performance with a mere 256 data samples (128 each for  $D^{\text{hon}}$  and  $D^{\text{task}}$ ) while modifying just 20% of the model’s parameters. Furthermore, as a training-free method, HCNR is inherently more efficient, delivering at least a 2.23x speedup over all baselines. These results confirm that HCNR offers a practical path to honesty recovery that is efficient in both data and time, making it viable for real-world deployment.

**More experiments based on other LLMs and FFT.** We also evaluate it on Llama-3-8B-Instruct, as well as Qwen and Mistral series models. Additionally, we conduct experiments using full fine-tuning (FFT) to further assess HCNR across training paradigms. The results are shown in Appendix D, demonstrating that HCNR consistently achieves honesty improvements and high efficiency across all these settings.

## Ablation Study

**Ablation Studies on HCNR’s Components.** To validate the effectiveness of our proposed HCNR framework, we conduct ablation studies on both core components. For Stage 1 (honesty-critical neurons recognition), we compare our method against Random selection (randomly select neurons) and w/o Task variants (using  $r_{j,k} = s_{j,k}^{\text{hon}}$  without task considerations). The number of neurons selected is the same across all experimental settings. Table 3 shows HCNR achieves F1 improvements over Random selection, while




Figure 6: Ablation study on the required dataset sizes  $|D^{task}|$  and  $|D^{hon}|$  for HCNR. We vary the dataset sizes and record the recovered models’ average F1 scores on UnknownBench and in-domain task accuracy.

the w/o Task variant causes domain task performance drops, demonstrating the necessity of critical neuron identification. For Stage 2 (honesty compensation), comparison with the w/o Compensation (w/o Com) variant reveals the critical importance of honesty compensation. When both stage components are simultaneously ablated (Random and w/o Com), honesty recovery becomes substantially impaired. These results validate both the precision of our neuron identification and the necessity of a compensation mechanism for effective honesty enhancement while preserving task-specific capabilities.

**Influence of  $|D^{task}|$  and  $|D^{hon}|$ .** To evaluate the data efficiency of the HCNR framework, we conduct ablation studies on the sizes of the two required datasets:  $D^{hon}$  and  $D^{task}$ . We systematically vary these dataset sizes while recovering HotpotQA fine-tuned models under the same settings. As Figure 6 shows, both honesty restoration and downstream accuracy plateau with only  $\approx 128$  examples. These findings confirm that the HCNR framework achieves optimal recovery performance for both honesty restoration and task preservation using remarkably few examples, highlighting its practical efficiency.

**Influence of  $R_{IW}$  and  $R_{CW}$ .** To investigate the impact of  $R_{IW}$  and  $R_{CW}$  hyperparameters on honesty recovery, we conduct ablation studies on honesty-impaired models, evaluating performance on KUQ and SelfAware datasets. As shown in Figure 7, both F1 score and Refusal  $\Delta$  improve progressively with increasing hyperparameter values. However,  $R_{IW}$  exhibits rapid saturation with diminishing marginal gains, while  $R_{CW}$  achieves optimal performance around 0.3, with further increases causing performance degradation.

## Related Works

### LLMs’ Honesty and Self-Knowledge

Honesty-enhancing methods can be categorized into two approaches: refusal-aware instruction tuning (RAIT) and reinforcement learning (Li et al. 2024). These approaches all require constructing large-scale, meticulously designed training datasets followed by retraining. RAIT methods (Zhang et al. 2023) evaluate LLMs’ knowledge boundaries by sampling and modifying labels for subsequent SFT. Reinforcement learning approaches (Cheng et al. 2024) enhance hon-

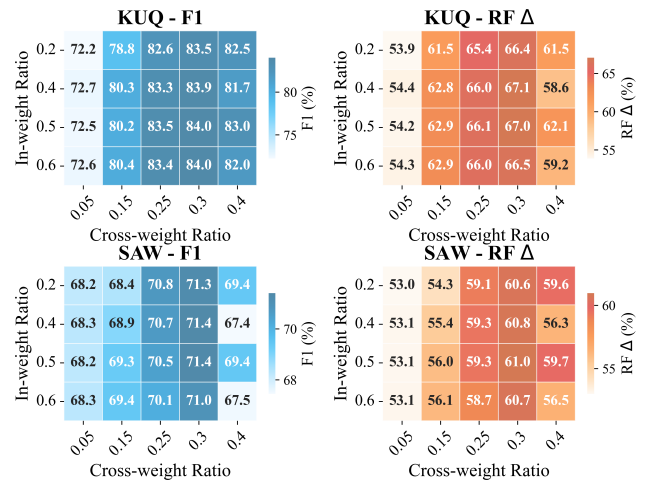


Figure 7: Ablation study of in-weight ratio  $R_{IW}$  and cross-weight ratio  $R_{CW}$  hyperparameters on honesty recovery for the fine-tuned model on the HotpotQA dataset.

esty through constructed preference data and reward modeling. Both approaches inevitably cause catastrophic forgetting and performance degradation on in-domain tasks. Moreover, all existing approaches overlook the underlying mechanisms of honesty deterioration, focusing solely on post-hoc correction.

## Knowledge and Capability Neurons

The concept of knowledge and capability neurons has emerged as a promising approach for attributing and localizing model behavior (Niu et al. 2024; Ferrando et al. 2024). Existing research has demonstrated that specific neurons correlate with various knowledge concepts (Dai et al. 2021; Shi et al. 2024). Recent advances have extended this line of inquiry to critical capabilities, revealing associations between specific neurons and safety (Yi et al. 2025), and confidence regulation (Stolfo et al. 2024). Building on the demonstrated success of neuron-level analysis across various domains, this work investigates honesty-related neurons to address the gap in understanding how large language models process honesty at the neuronal level.

## Conclusion

This work addresses honesty degradation in fine-tuned LLMs that leads to confident fabrications. Mechanistically, we reveal that dishonesty is spurious: while SFT alters expression behavior, inner knowledge boundary representations remain stable. This insight enables our Honesty-Critical Neurons Restoration (HCNR), which restores truthful behavior by selectively reverting honesty-critical neurons to pre-trained states and employing Hessian-guided compensation. Extensive experiments demonstrate that HCNR achieves substantial honesty recovery with minimal overhead while preserving task performance, providing an efficient solution for trustworthy AI in high-stakes domains.

## Acknowledgements

The work is supported by the grants from the Natural Science Foundation of China (62225202, 62202029), and Young Elite Scientists Sponsorship Program by CAST (No. 2023QNRC001). We owe sincere thanks to all authors for their valuable efforts and contributions. The corresponding author is Tianyu Chen and Haoyi Zhou.

## References

- Amayuelas, A.; Wong, K.; Pan, L.; Chen, W.; and Wang, W. 2023. Knowledge of knowledge: Exploring known-unknowns uncertainty with large language models. *arXiv preprint arXiv:2305.13712*.
- Askell, A.; Bai, Y.; Chen, A.; Drain, D.; Ganguli, D.; Henighan, T.; Jones, A.; Joseph, N.; Mann, B.; and Das-sarma, N. 2021. A General Language Assistant as a Laboratory for Alignment.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Cheng, Q.; Sun, T.; Liu, X.; Zhang, W.; Yin, Z.; Li, S.; Li, L.; He, Z.; Chen, K.; and Qiu, X. 2024. Can AI assistants know what they don't know? *arXiv preprint arXiv:2401.13275*.
- Dahl, M.; Magesh, V.; Suzgun, M.; and Ho, D. E. 2024. Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*, 16(1): 64–93.
- Dai, D.; Dong, L.; Hao, Y.; Sui, Z.; Chang, B.; and Wei, F. 2021. Knowledge neurons in pretrained transformers. *arXiv preprint arXiv:2104.08696*.
- Ferrando, J.; Sarti, G.; Bisazza, A.; and Costa-Jussà, M. R. 2024. A primer on the inner workings of transformer-based language models. *arXiv preprint arXiv:2405.00208*.
- Glaese, A.; McAleese, N.; Trébacz, M.; Aslanides, J.; Firoiu, V.; Ewalds, T.; Rauh, M.; Weidinger, L.; Chadwick, M.; Thacker, P.; et al. 2022. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*.
- Hassibi, B.; Stork, D. G.; and Wolff, G. J. 1993. Optimal brain surgeon and general network pruning. In *IEEE international conference on neural networks*, 293–299. IEEE.
- Hong, J.; Lee, N.; and Thorne, J. 2024. Orpo: Monolithic preference optimization without reference model. *arXiv preprint arXiv:2403.07691*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Kim, Y.; Jeong, H.; Chen, S.; Li, S. S.; Lu, M.; Alhamoud, K.; Mun, J.; Grau, C.; Jung, M.; Gameiro, R.; et al. 2025. Medical hallucinations in foundation models and their impact on healthcare. *arXiv preprint arXiv:2503.05777*.
- Krithara, A.; Nentidis, A.; Bougiatiotis, K.; and Paliouras, G. 2023. BioASQ-QA: A manually curated corpus for Biomedical Question Answering. *Scientific Data*, 10(1): 170.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7: 453–466.
- LeCun, Y.; Denker, J.; and Solla, S. 1989. Optimal brain damage. *Advances in neural information processing systems*, 2.
- Li, S.; Yang, C.; Wu, T.; Shi, C.; Zhang, Y.; Zhu, X.; Cheng, Z.; Cai, D.; Yu, M.; Liu, L.; et al. 2024. A survey on the honesty of large language models. *arXiv preprint arXiv:2409.18786*.
- Liu, G.; Wang, X.; Yuan, L.; Chen, Y.; and Peng, H. 2023. Examining LLMs' Uncertainty Expression Towards Questions Outside Parametric Knowledge. *arXiv preprint arXiv:2311.09731*.
- Mai, Z.; Chowdhury, A.; Zhang, P.; Tu, C.-H.; Chen, H.-Y.; Pahuja, V.; Berger-Wolf, T.; Gao, S.; Stewart, C.; Su, Y.; et al. 2024. Fine-tuning is fine, if calibrated. *arXiv preprint arXiv:2409.16223*.
- Manish, S. 2023. Constitutional AI: An Expanded Overview of Anthropic's Alignment Approach. *Information Horizons: American Journal of Library and Information Science Innovation*, 1(7): 36–39.
- Nguyen, H.; He, Z.; Gandre, S. A.; Pasupulety, U.; Shivakumar, S. K.; and Lerman, K. 2025. Smoothing Out Hallucinations: Mitigating LLM Hallucination with Smoothed Knowledge Distillation. *arXiv preprint arXiv:2502.11306*.
- Niu, J.; Liu, A.; Zhu, Z.; and Penn, G. 2024. What does the Knowledge Neuron Thesis Have to do with Knowledge? *arXiv preprint arXiv:2405.02421*.
- Pal, A.; Umapathi, L. K.; and Sankarasubbu, M. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, 248–260. PMLR.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741.
- Sarabadani, S. 2019. Detection of adverse drug reaction mentions in tweets using ELMo. In *Proceedings of the fourth social media mining for health applications (#smm4h) workshop & shared task*, 120–122.
- Shi, D.; Jin, R.; Shen, T.; Dong, W.; Wu, X.; and Xiong, D. 2024. Ircan: Mitigating knowledge conflicts in llm generation via identifying and reweighting context-aware neurons. *Advances in Neural Information Processing Systems*, 37: 4997–5024.
- Stolfo, A.; Wu, B.; Gurnee, W.; Belinkov, Y.; Song, X.; Sachan, M.; and Nanda, N. 2024. Confidence regulation neurons in language models. *Advances in Neural Information Processing Systems*, 37: 125019–125049.
- Wang, Y.; Bai, A.; Peng, N.; and Hsieh, C.-J. 2024. On the loss of context-awareness in general instruction fine-tuning. *arXiv preprint arXiv:2411.02688*.

Xu, T.; Wu, S.; Diao, S.; Liu, X.; Wang, X.; Chen, Y.; and Gao, J. 2024. Saysself: Teaching llms to express confidence with self-reflective rationales. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 5985–5998.

Yang, Y.; Chern, E.; Qiu, X.; Neubig, G.; and Liu, P. 2024. Alignment for honesty. *Advances in Neural Information Processing Systems*, 37: 63565–63598.

Yang, Z.; Qi, P.; Zhang, S.; Bengio, Y.; Cohen, W. W.; Salakhutdinov, R.; and Manning, C. D. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.

Yi, X.; Zheng, S.; Wang, L.; de Melo, G.; Wang, X.; and He, L. 2025. Nlsr: Neuron-level safety realignment of large language models against harmful fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 25706–25714.

Yin, Z.; Sun, Q.; Guo, Q.; Wu, J.; Qiu, X.; and Huang, X. 2023. Do large language models know what they don't know? *arXiv preprint arXiv:2305.18153*.

Zhang, H.; Diao, S.; Lin, Y.; Fung, Y. R.; Lian, Q.; Wang, X.; Chen, Y.; Ji, H.; and Zhang, T. 2023. R-tuning: Teaching large language models to refuse unknown questions. *arXiv preprint arXiv:2311.09677*.

Zhang, X.; and Wu, J. 2024. Dissecting learning and forgetting in language model finetuning. In *The Twelfth International Conference on Learning Representations*.

Zhang, Z.; Wang, F.; Li, X.; Wu, Z.; Tang, X.; Liu, H.; He, Q.; Yin, W.; and Wang, S. 2024. Catastrophic Failure of LLM Unlearning via Quantization. *arXiv preprint arXiv:2410.16454*.