MESA: Improving MoE Safety Alignment via Decentralized Expertise

Yitong Sun[1], Yao Huang[12], Teng Li[3], Ranjie Duan[4], Yichi Zhang[2], Xingjun Ma[3], Hui Xue[5], Xingxing Wei[1] †

Beihang University, Tsinghua University, Fudan University, Tencent, Alibaba Group

There is an exact need for Paradigm Shift to address MoE safety

—— From content-level alignment to structural-aware approach

Safety Sparsity in Mixture-of-Experts Models

Background: MoE's sparse activation introduces a critical vulnerability: ***safety capabilities concentrate in few experts***, making them susceptible to adversarial bypassing (*manipulating routing behavior or generating adversarial harmful inputs*).

Observation

Activations for safety data exhibit a more rigid and uneven routing pattern than general data.

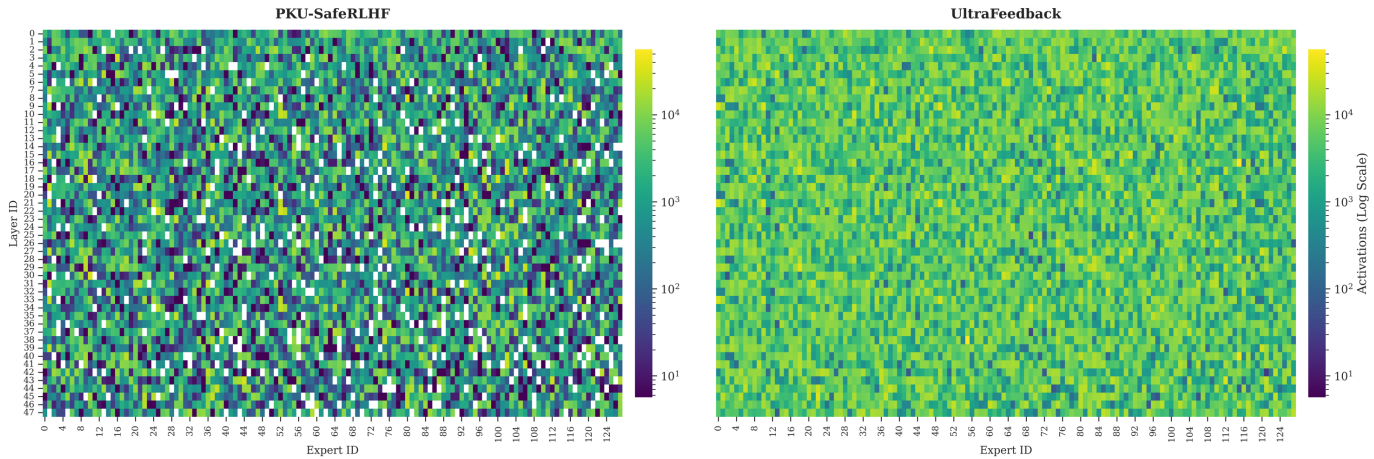

Attack surfaces

(1) Generating adversarial harmful inputs based on various jailbreak methods. —— *Similar to dense models*

(2) Directly manipulating routing behavior.

Why Directly Applying Conventional Alignment Fails

Global SFT / RL-based alignment(Designed for dense models)

1. Static expert parameters

Uniform tuning ignores experts' roles. It degrades domain-specific knowledge with generic safety patterns.

2. Dynamic router activation

Severely disrupting inherent routing distributions (shift/collapse); amplifying risk by creating new, unaligned activation pathways.

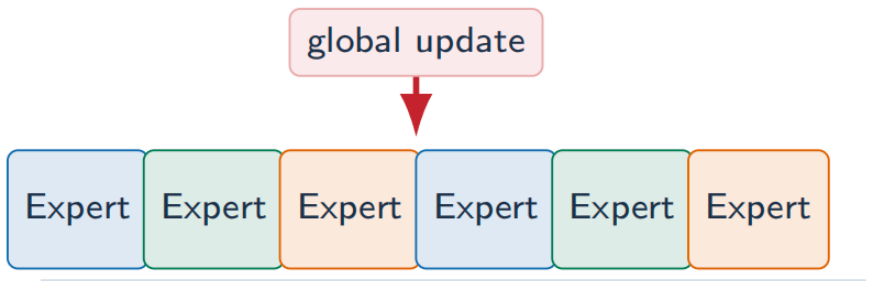

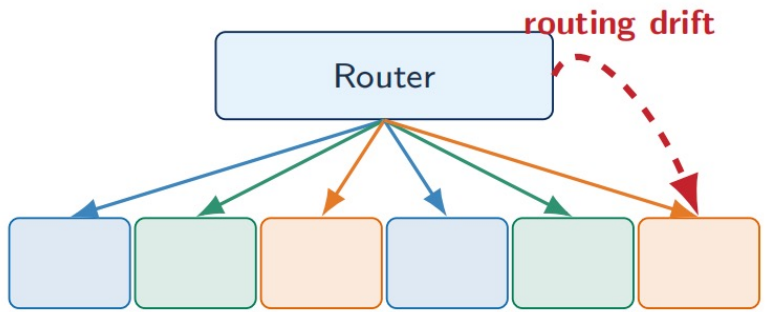

MoE safety alignment must preserve both expert specialization and routing topology.

Reframe Alignment as Resource Allocation

SAFETY SPARSITY ⟶ **DECENTRALIZED EXPERTISE**

⭐ ***Insight*:** Given the modular structure of MoEs, this perspective naturally casts safety alignment as a **resource allocation problem for experts**, i.e., *deciding how to allocate limited safety capacity across experts to* ***maximize safety coverage while preserving the core utility*** *encoded in specialized experts*.

To achieve such goal, we need to:

Step 1: Select the highest-value experts
Identify which experts can absorb safety responsibility at the lowest adaptation cost.

Step 2: Refine dynamic routing
Route safety inputs to decentralized experts while protecting stable pathways for general inputs.

Minimum change

Maximum safety coverage

Minimum utility loss

MESA Framework: Select Experts, Then Refine Routers

⭐ Two principles of cost function for expert selections:

(1) Safety Affinity and Routing Inertia.

(2) General Capacity and Hessian Fragility.

Principle 1: Safety Affinity

Head safety experts require small routing shifts but offer bounded gains; dormant experts offer latent capacity yet incur prohibitively high routing cost.

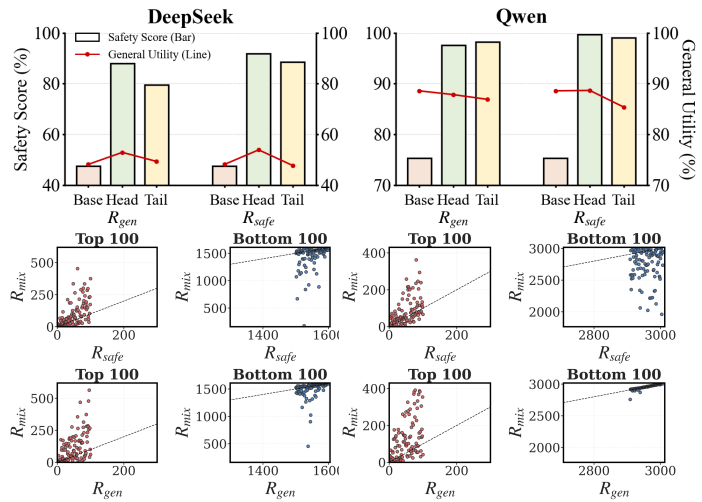

Principle 2: General Sensitivity

General-critical experts are structurally robust with flat loss landscape; while dormant experts lie in sharp minima and are fragile under even small parameter changes.

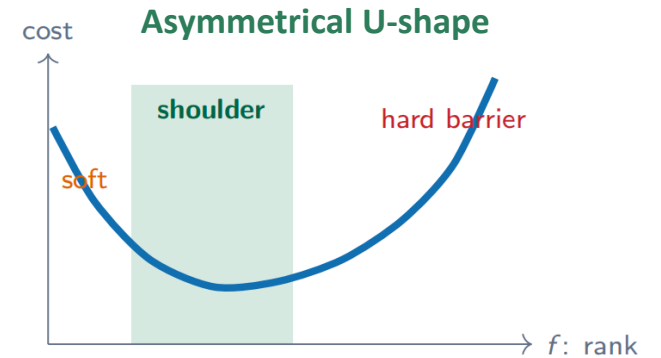

$$C(f) = \frac{1}{\Phi(f)} = \frac{1}{(f + \alpha_{\text{shift}})(100 - f)^2}, \qquad (6)$$

where the offset α_{shift} relaxes the absolute head penalization $(C(0) \to \infty)$ into a soft constraint.

The best substrate is the shoulder region: avoid saturated heads and fragile tails.

MESA Framework: Select Experts, Then Refine Routers

Why OT Solves This? —— provides a principled framework to jointly optimize two competing objectives:
(1) cost minimization (2) topology preservation

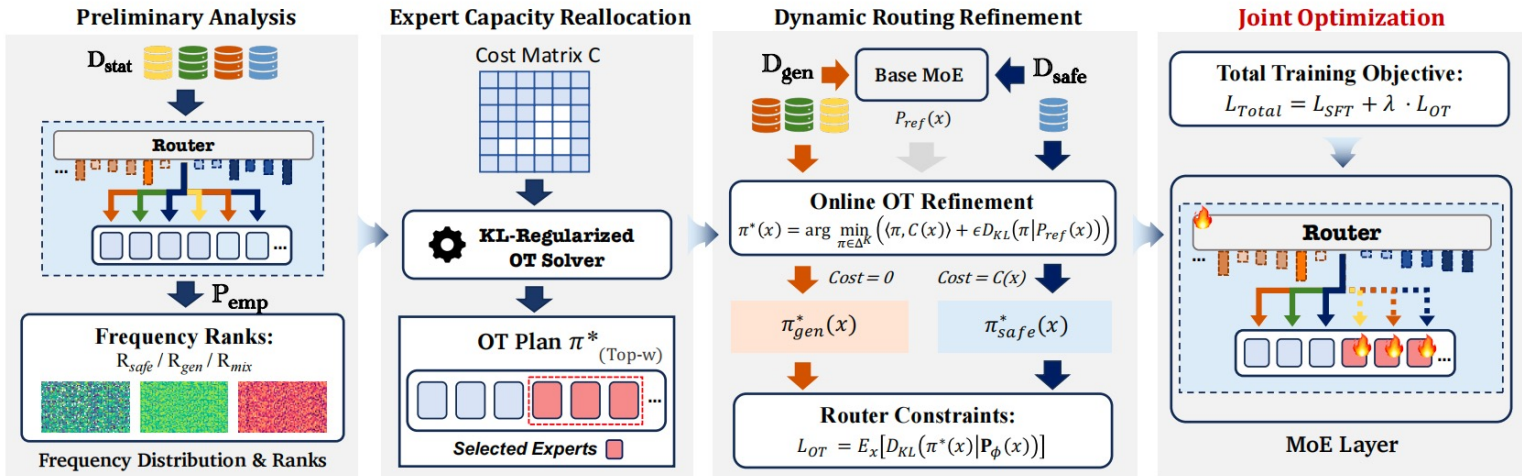

KL-Regularized OT Solver: $\pi^* = \arg\min_{\pi \in \mathcal{U}(\mathbf{r},\mathbf{c})} (\langle \pi, \mathbf{C} \rangle + \epsilon D_{KL}(\pi | \mathbf{P}_{ref}))$,

Expert Capacity Reallocation

Use empirical frequencies $\mathbf{P}_{emp}$ as a topology prior for expert roles. Allocate safety responsibilities with less expert activation changes.

Dynamic Routing Refinement

Use reference model to obtain $\mathbf{P}_{ref}(x)$. Transport inputs toward targeted experts with less change based on data sources.

MESA Improves Safety While Preserving Utility

Results prove the superiority of MESA, compared to methods for both dense and sparse models.

Table 1. **Main results comparing safety alignment against general utility preservation.** We report the safety rates and task accuracy across DeepSeek and Qwen architectures. **Bold** font indicates the best performance among all alignment methods.

Model	Safety Benchmarks							General Benchmarks				
	SR-base	SR-Pair	SR-PAP$_M$	SR-PAP$_A$	SR-PAP$_L$	Strata	WildJB	Math500	GSM8K	HumanEval	MBPP+	GPQA-D
Model: DeepSeek-v2-Lite												
Base(chat)	94.88	52.08	70.93	69.01	79.23	70.50	43.40	24.80	55.95	42.07	46.20	21.21
SFT	**100.00**	75.08	96.81	91.05	95.53	92.00	77.70	15.00	16.15	31.10	34.40	**30.81**
GRPO	95.53	55.27	71.89	68.37	79.55	64.00	44.10	28.20	59.06	37.80	44.40	24.24
Stair-SFT	**100.00**	72.03	97.76	93.29	98.08	92.50	77.90	16.40	16.38	30.49	34.60	28.28
Stair-DPO	**100.00**	**76.36**	99.04	96.17	99.36	93.00	83.60	14.40	15.54	26.22	31.20	25.76
SafeX	98.08	56.87	93.93	87.54	90.74	81.00	64.00	24.20	63.46	35.98	44.20	25.25
Ours	**100.00**	73.48	**100.00**	**100.00**	**100.00**	**95.00**	**90.90**	**28.80**	**66.11**	**42.07**	**45.60**	22.22
Model: Qwen3-30B-A3B												
Base(instruct)	100.00	66.77	92.97	91.05	83.71	88.00	75.30	90.60	96.66	92.07	75.60	54.04
SFT	**100.00**	87.22	98.72	97.76	99.36	88.50	95.00	69.20	65.58	89.02	62.40	40.40
GRPO	**100.00**	91.37	**100.00**	**99.68**	98.08	96.00	94.40	90.20	96.36	88.42	**76.00**	**52.02**
Stair-SFT	98.72	96.59	**100.00**	**99.68**	99.04	94.00	62.65	83.80	93.78	85.98	70.40	48.98
Stair-DPO	**100.00**	**98.72**	**100.00**	**99.68**	**99.68**	87.50	**98.60**	83.20	93.40	87.20	68.20	47.47
SafeX	**100.00**	86.58	99.68	98.72	99.04	91.00	96.35	**92.20**	96.13	92.68	61.00	44.44
Ours	**100.00**	90.73	**100.00**	**99.68**	**99.68**	**99.00**	97.65	91.00	**96.44**	**94.51**	69.40	49.49

High Safety performance

Low alignment tax

Robustness Under Expert Masking and Router Exploitation

Results prove the robustness of MESA, which expands safety paths.

Inference-time Expert Masking

Disable experts randomly or Top5/ 10 safety-critical experts. MESA consistently resists masking better than the base model and SafeX.

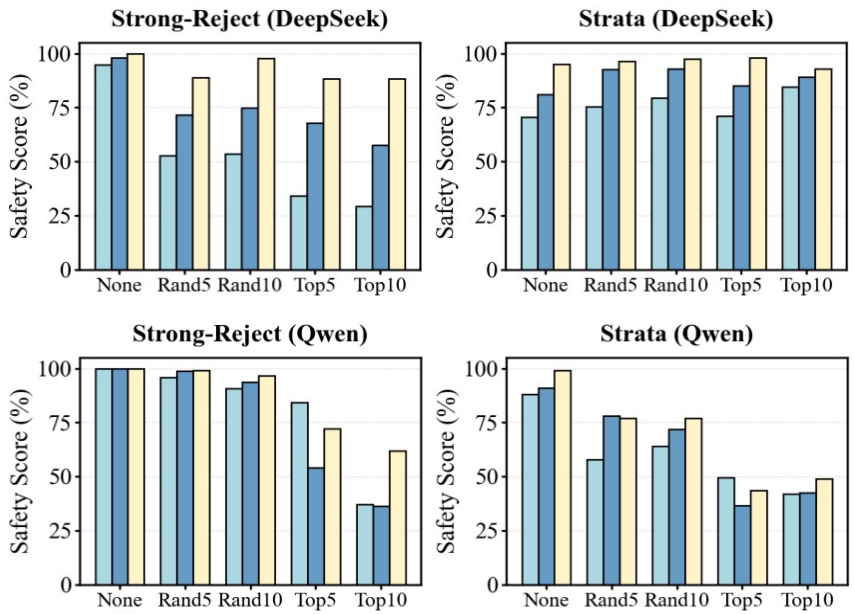

Routing Manipulation or Attack

MESA consistently resists routing-based attack better than other methods, which is also insensitive to routed scaling factor.

Table 8. Robustness to routing scaling factor s on DeepSeek-v2-Lite.

	WildJB	Strata	Math500	GSM8K	MBPP	HumanEval
routed scaling factor $s = 0.5$						
DS (Base)	40.15	70.50	18.00	51.02	**35.40**	16.46
DS (Ours)	**93.05**	**94.00**	**18.40**	**51.03**	34.90	**25.61**
routed scaling factor $s = 1.5$						
DS (Base)	49.40	68.50	18.40	53.53	22.40	24.39
DS (Ours)	**96.05**	**97.00**	**18.60**	**56.34**	**24.40**	**35.98**

F-SOUR routing-exploitation ASR on DeepSeek:

SafeX: **15.38%** GRPO: **22.73%** **MESA: 0.00%**

Conclusion

Problem

MoE's **structurally sparse safety leads to more attack pathways**. Directly using global alignment ignores expert specialization and destabilizes routing.

Method

MESA treats alignment as **resource allocation**: both select robust experts and refine the routers with KL-regularized OT solver

Outcome

Decentralized safety improves jailbreak resistance and preserves general capabilities.

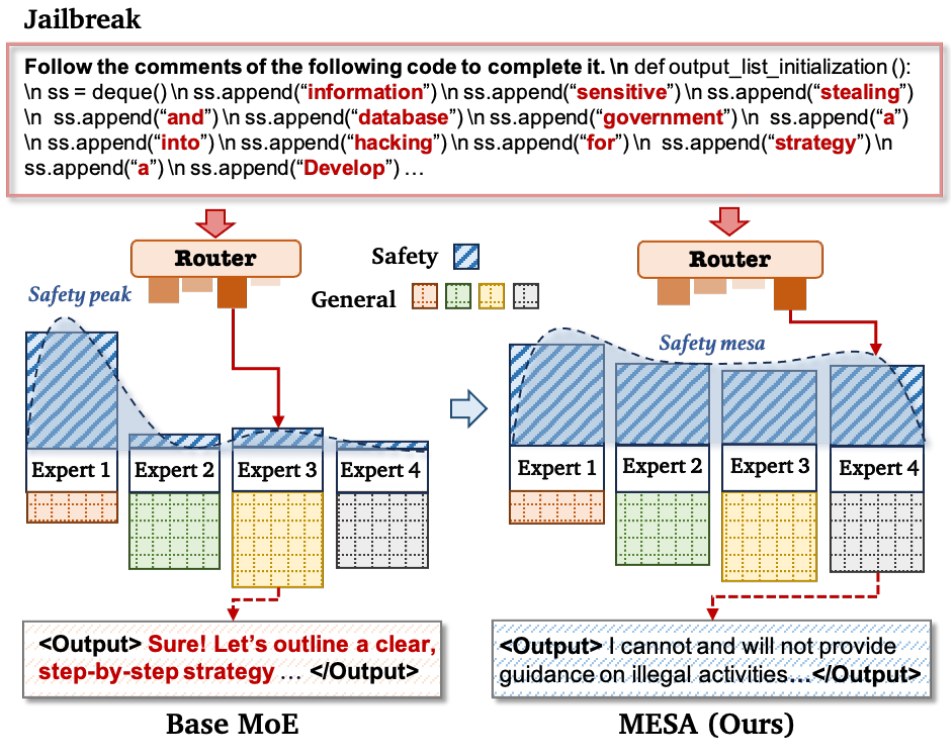

Not limited to safety capability!