PEFT-Arena: Understanding Parameter-Efficient Finetuning from a Stability-Plasticity Perspective

Yangyi Huang 1155251471 yyhuang@link.cuhk.edu.hk
Department of Computer Science and Engineering, Faculty of Engineering, CUHK

Research Questions

RQ1

Which PEFT method achieves the best stability-plasticity trade-off?

Benchmark PEFT by **target adaptation** and **general capability retention** across math and medical reasoning.

RQ2

What internal mechanisms govern the stability-plasticity trade-off?

Probe PEFT through **spectral geometry**, then use **interpolation** to turn that understanding into better operating points.

Overview

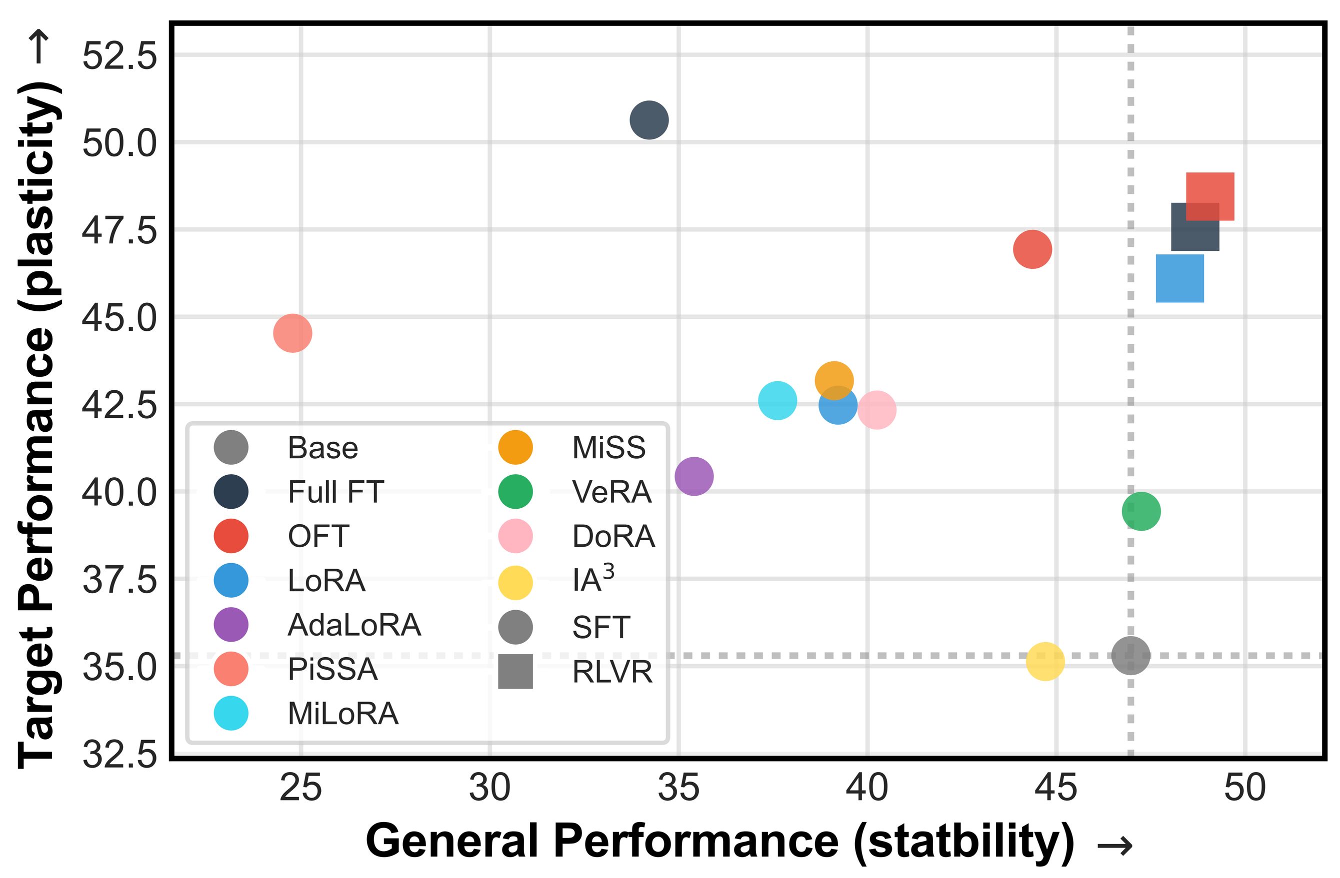

(a) Results on PEFT-Arena

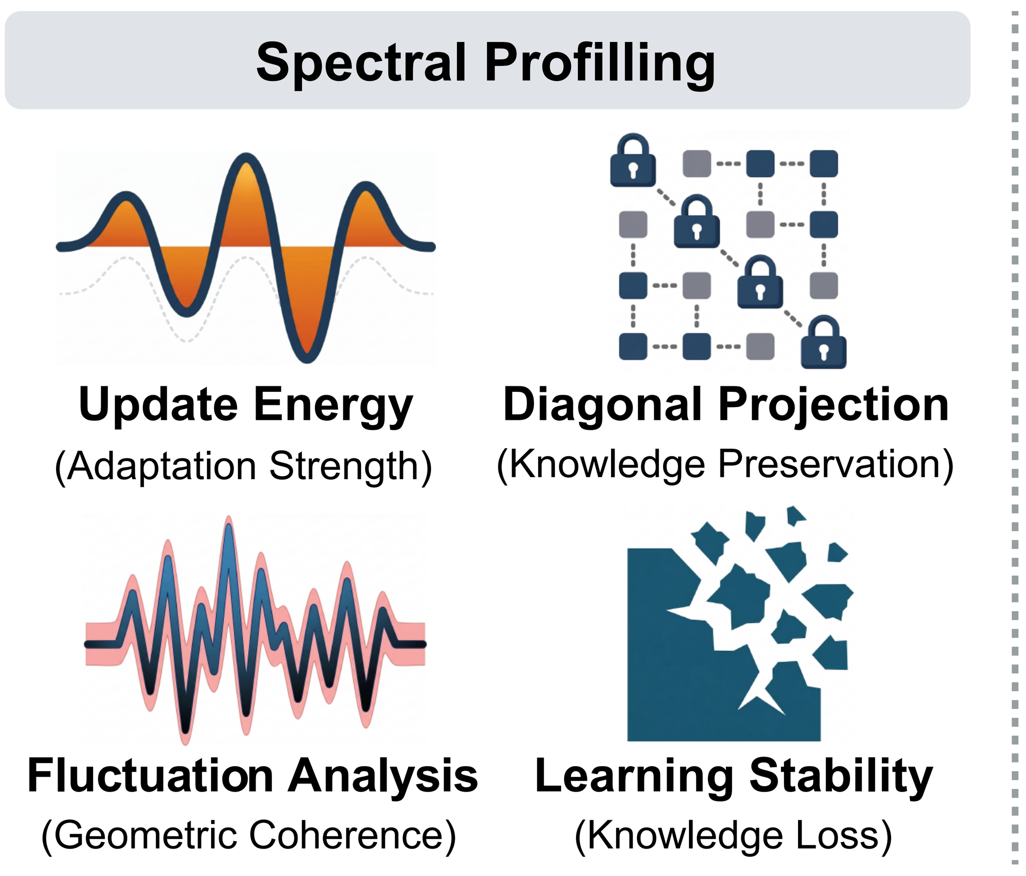

(b) Internal Spectral Analysis

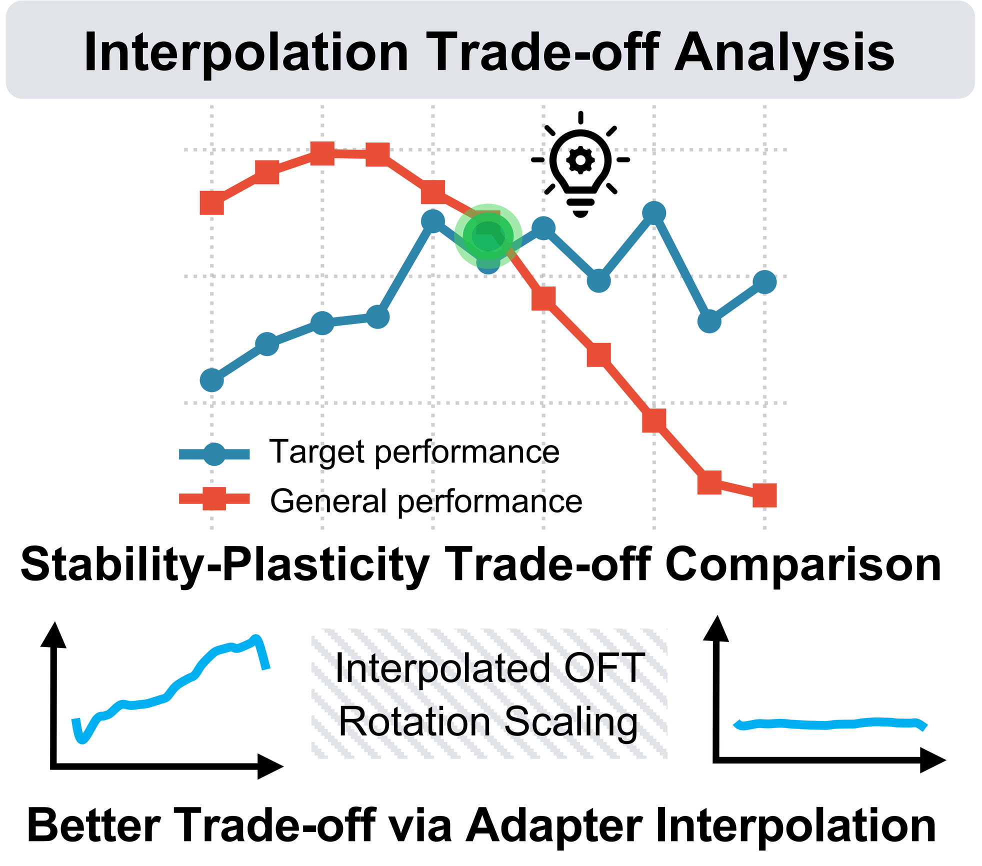

(c) External Performance Analysis

Part I. Benchmark
Evaluate target gains together with retained general capability.

Part II. Internal Mechanism
Spectral profiling reveals why some updates forget less.

Part III. Trade-off by Interpolation
Overshoot motivates interpolation and the geometry-aware iOFT.

Part I. PEFT-Arena Benchmark

We evaluate both **Qwen2.5-7B-base** and **Llama3.2-3B-Instruct** under supervised fine-tuning (**SFT**) and **GRPO-based reinforcement learning**, measuring both **target** and **general** performance.

Method	H.P.	Q Tr.	Qwen2.5-7B-base Math T	Math G	Med T	Med G	L Tr.	Llama3.2-3B-Instruct Math T	Math G	Med T	Med G
SFT											
Base	–	7.61B	35.30	46.97	46.36	46.97	3.21B	27.63	53.03	41.44	56.76
Full FT	–	6.53B	50.63	34.22	53.63	34.41	2.85B	33.90	39.83	44.26	26.03
OFT	b16	8.49M	42.33	42.58	46.17	45.09	7.08M	29.43	41.08	39.22	40.97
OFT	b32	17.55M	46.93	44.37	48.63	42.40	11.55M	30.60	40.73	39.50	40.50
OFT	b64	35.68M	46.23	35.97	49.47	39.11	24.97M	29.30	39.75	40.77	37.70
OFT	b128	71.92M	47.77	36.98	52.40	36.88	47.34M	32.23	36.26	42.17	34.26
LoRA	r4a8	10.09M	42.33	41.66	47.12	36.42	6.90M	24.30	35.79	36.92	31.84
LoRA	r8a16	20.19M	42.47	39.22	47.91	36.06	12.16M	24.07	36.57	38.34	27.99
LoRA	r16a32	40.37M	44.87	34.91	47.86	34.86	24.31M	24.97	37.55	39.21	29.18
LoRA	r32a64	80.74M	45.37	38.21	49.48	35.50	48.63M	25.90	37.20	39.33	30.69
AdaLoRA	r8a16	30.28M	40.43	35.41	45.22	37.34	18.24M	20.83	34.53	37.11	36.29
PiSSA	r8a16	20.19M	44.53	24.78	26.16	18.05	12.16M	0.67	7.08	21.17	9.75
MiLoRA	r8a16	20.19M	42.60	37.62	46.83	35.88	12.16M	23.60	35.59	37.64	29.23
MiSS	r8	11.12M	43.17	39.12	48.75	34.43	6.19M	23.37	33.93	40.16	31.71
MiSS	r64	89.00M	46.93	32.77	51.90	32.72	49.55M	28.63	34.96	41.96	27.07
VeRA	r256	1.44M	39.43	47.25	28.50	47.01	0.82M	28.80	46.79	40.68	48.94
DoRA	r8a16	21.58M	42.33	40.25	48.04	36.06	12.93M	23.83	35.65	38.25	27.53
IA^3	–	1.82M	35.13	44.71	29.70	45.72	0.92M	29.70	45.72	39.13	45.67
RLVR with GRPO											
Full FT	–	6.53B	47.57	48.68	46.24	43.22	2.85B	29.80	52.20	45.88	51.81
OFT	b32	8.49M	48.37	48.90	46.79	47.24	11.55M	29.97	50.04	44.99	52.31
LoRA	r8a16	20.19M	46.10	48.27	47.08	42.80	24.31M	28.83	52.17	46.24	53.54

Takeaways.

- Full FT tends to learn more and forget more, while PEFT behavior depends on both parameter budget and parameterization.
- Within additive LoRA-style PEFT, SVD-space methods do not yield stable performance in SFT.
- OFT, based on orthogonal rotation, achieves the best **stability-plasticity trade-off**.

Part II. Spectral Mechanism

Key idea. Better retention aligns with **smoother spectral deviations**.

Retention profile (diagonal projection)
$P_{\mathrm{diag}}(i) = u_i^\top W_{\mathrm{ft}} v_i$

Adaptation profile (update energy)
$E_\Delta(i) = \|(W_{\mathrm{ft}} - W_{\mathrm{pre}}) v_i\|_2$

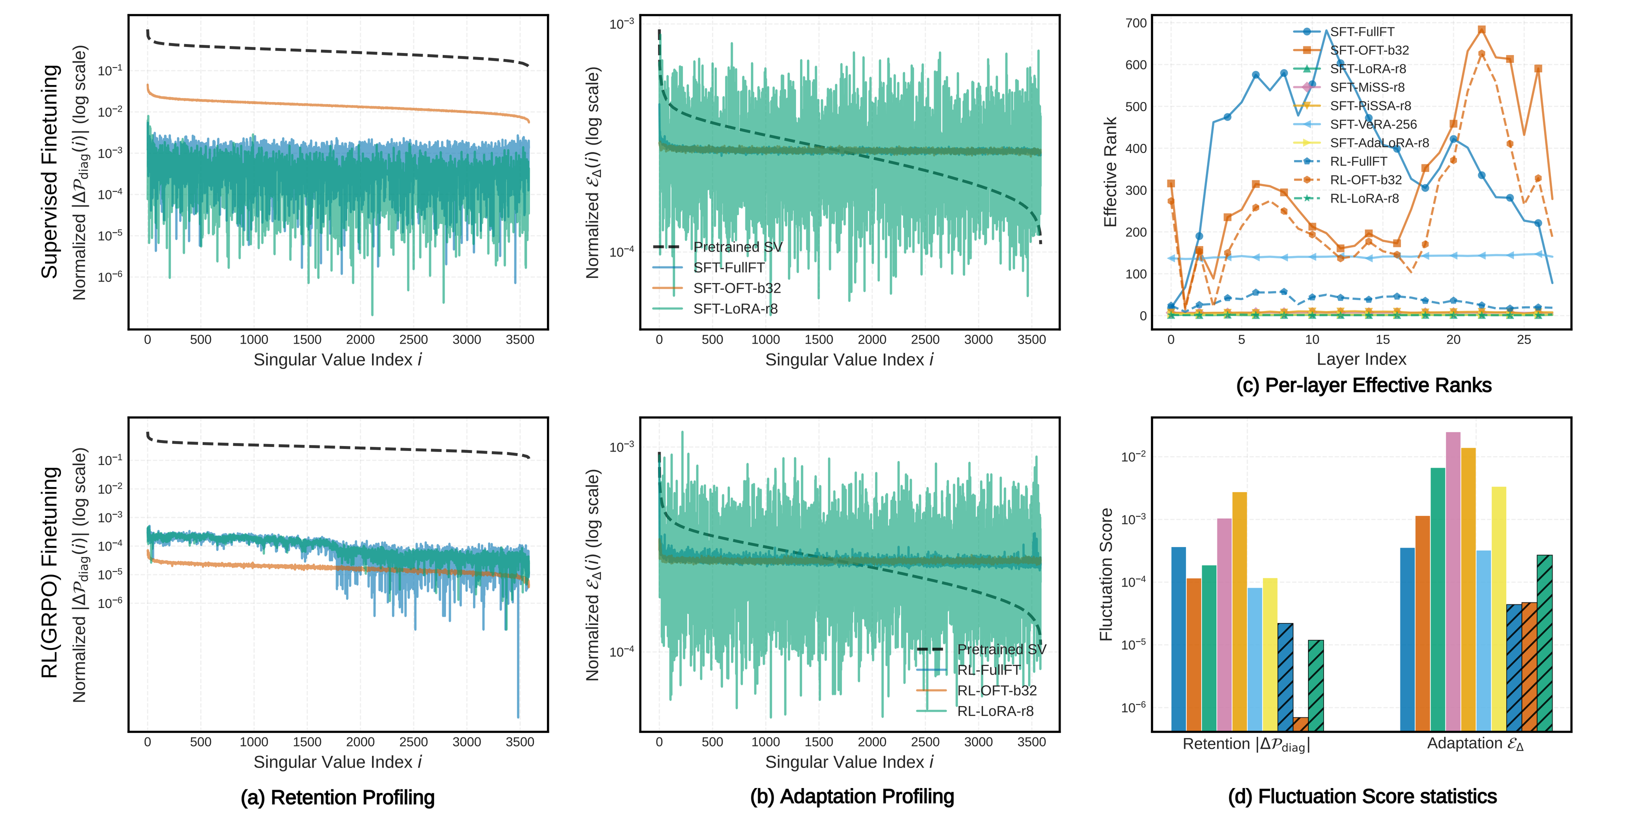

What the profiling shows.

- Additive low-rank PEFT is often spectrally spiky under SFT.
- OFT stays in a smoother, more globally systematic spectral regime via near-orthogonal rotations.
- Smoother retention and adaptation profiles explain why OFT forgets less while staying adaptive.

Part III. Interpolation, Overshoot, and iOFT

Trade-off curves. Interpolating between the base and finetuned checkpoints consistently improves the frontier.

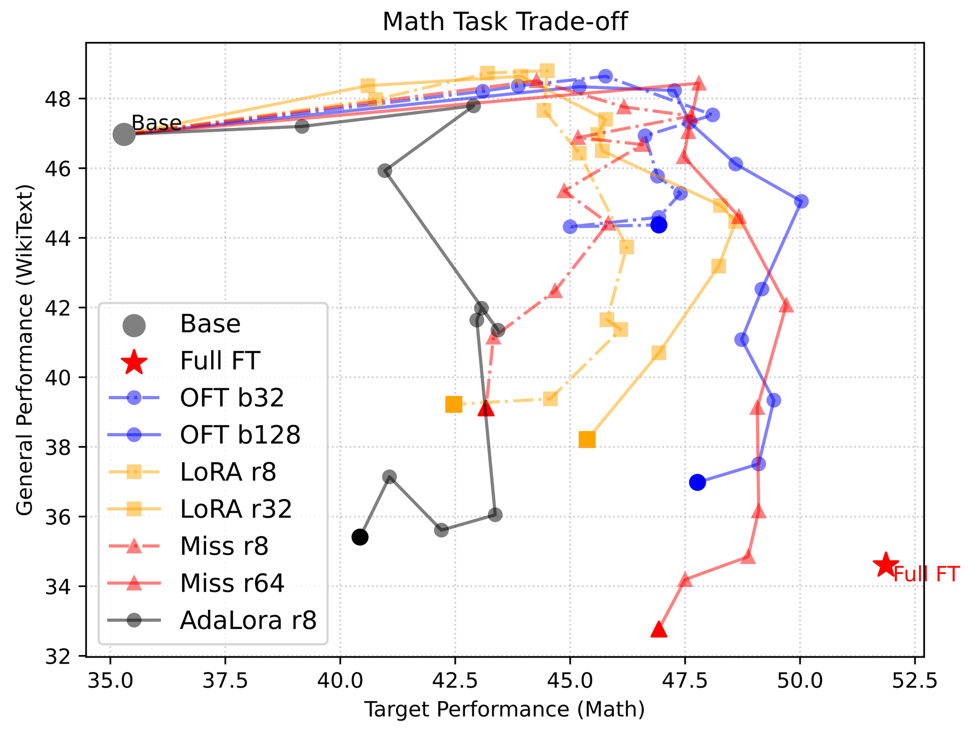

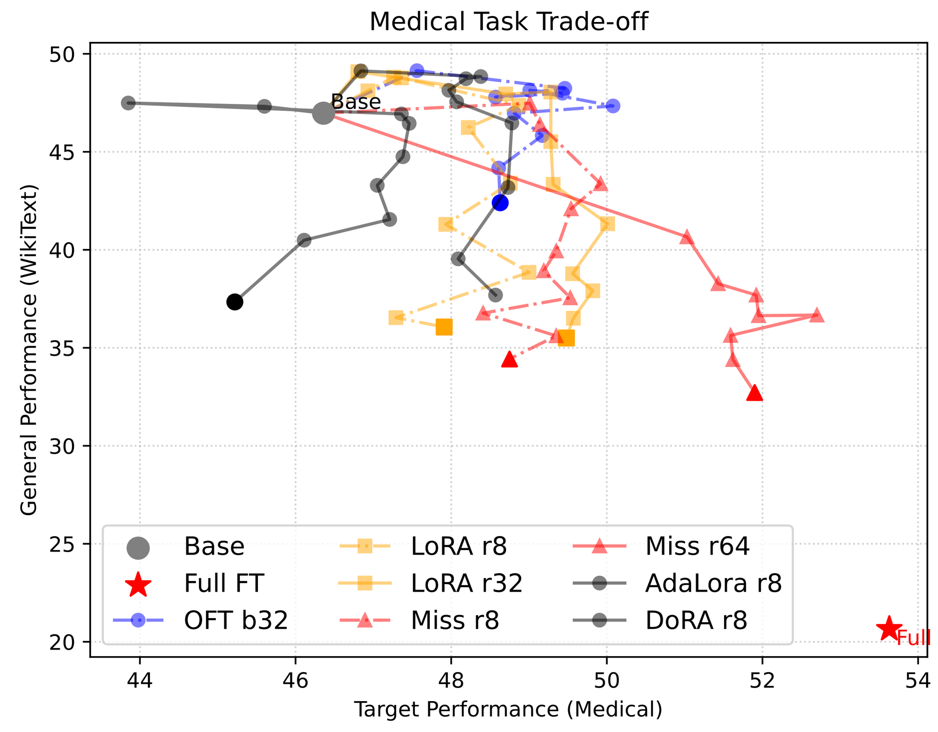

The Overshoot Phenomenon. Misalignment between training trajectories and parameter interpolation trajectories inspired our trade-off method.

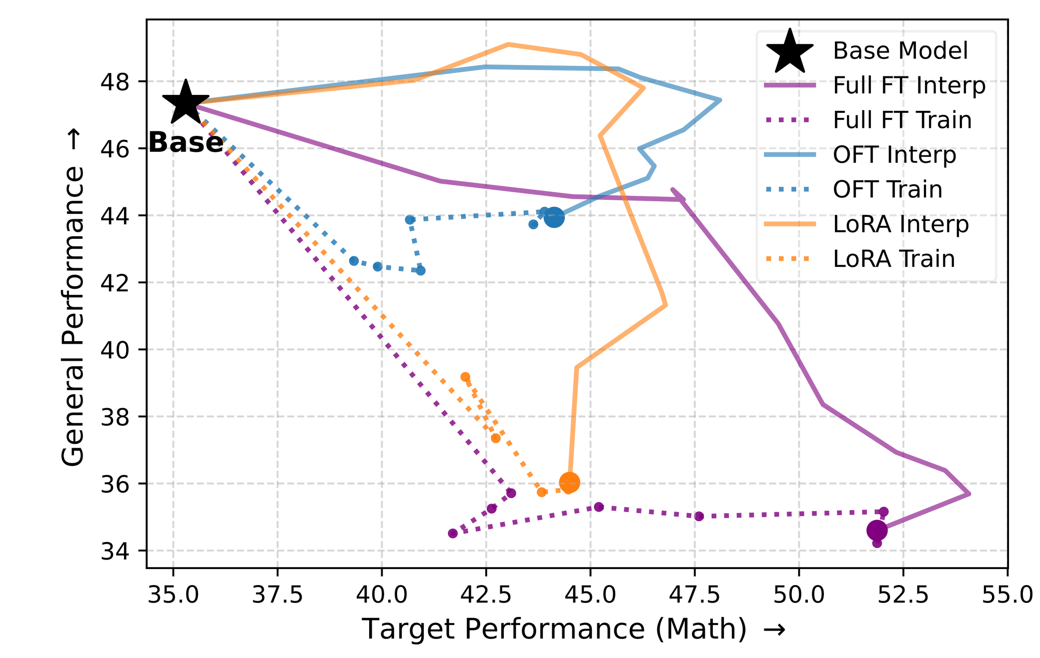

iOFT: geometry-aware OFT adjustment.

- Rescale Q_ℓ layer by layer to flatten rotation strength.
- **SafeRho**: match the first-five-layer average.
- **MinRho**: match the minimum across layers.
- Recover general ability without sweeping a global interpolation coefficient.

Variant	MT	MG	MedT	MedG
OFT	46.93	44.37	48.63	42.40
SafeRho	47.17	46.69	50.01	47.61
MinRho	47.83	46.86	49.76	47.79

Method	Δpass@1	Δpass@64	Δ(p@64 - p@1)
LoRA	2.30 ↑	2.93 ↓	5.23 ↓
OFT	0.07 ↓	1.46 ↓	1.39 ↓

RLVR after longer GRPO. OFT keeps high-k behavior much better than LoRA (avoid policy collapse with better diversity).