



Viacheslav Sinii, Alexey Gorbатовski, Artem Cherepanov, Boris Shaposhnikov, Nikita Balagansky, Daniil Gavrilov. T-Tech, Central University

# Think small.

Training steering vectors instead of all model weights doesn't change reasoning quality\*, but it makes a huge difference in cost and **interpretability**.

Average Performance	Qwen2.5-Math		Qwen2.5			LLaMa3.1-8B	
	1.5B	7B	1.5B	7B	14B	Base	It
Base	18.1	24.8	0.3	24.2	26.1	3.9	21.7
Full-Tune	33.6	43.5	22.7	37.1	41.3	11.5	26.1
Steering	33.9	43.3	23.2	36.4	42.3	9.1	26.1

LLaMa3.1-8B-it		
Resources	Full	Steering
Optimizer State	7.5Gb	128Kb
Weight-Broadcast	5.32s	0.07s
Broadcast Overall	28m	22s



\*Trained on DeepScaleR with the RLOO objective; evaluated on AIME24/25, AMC23, MATH500, Minerva-Math, and OlympiadBench. Table shows average benchmark results.