

Training configuration: Cosmos3-Nano-Policy + 💡 ARC

Parameter	Value
Model and objective	
Base model	Cosmos3-Nano-Policy.
Released DROID checkpoint	nvidia/Cosmos3-Nano-Policy-DROID.
Architecture	Mixture-of-Transformers: an autoregressive reasoner and a diffusion generator.
Trainable components	Joint fine-tuning of the reasoner and generator.
Training conditioning	Current observation, task instruction, and ground-truth causal trace.
Generator objective	Separate velocity flow-matching losses for future observations and actions; action-loss coefficient 10.
Trace objective	Next-token prediction of the causal trace.
Total objective	$\mathcal{L}_{\text{WAM}} = \mathcal{L}_{\text{FM}}^{\text{obs}} + 10\mathcal{L}_{\text{FM}}^{\text{act}} + \lambda_{\text{trace}}\mathcal{L}_{\text{trace}}$.
Training configuration	
Sequence length	33.
Batch size per rank	64.
Resolution setting	480.
Learning rate	2×10^{-4} .
Run end iteration	25 450.
Hardware and measured runtime	
Compute nodes	32.
GPUs	128 NVIDIA GB200 GPUs.
Latest iteration time	22.09 seconds per iteration.
Measured run wall time	154.29 hours.