Robotics

Required

Black et al., π_0: A Vision-Language-Action Flow Model for General Robot Control · arXiv:2410.24164

Optional

Brohan et al., RT-2 · arXiv:2307.15818 Kim et al., OpenVLA · arXiv:2406.09246

Chi et al., Diffusion Policy · arXiv:2303.04137 Zhao et al., ACT · arXiv:2304.13705

Hu et al., GAIA-1 · arXiv:2309.17080 NVIDIA, Cosmos World Foundation Model Platform · arXiv:2501.03575

Christoffer Heckman
Department of Computer Science, University of Colorado Boulder

November 19, 2026

Quick overview

Actions as tokens: RT-2 and OpenVLA

Flow matching and action chunks

π_0 to $\pi_{0.7}$

Cosmos Reason, Alpamayo and GR00T N1

World models and GR00T 2

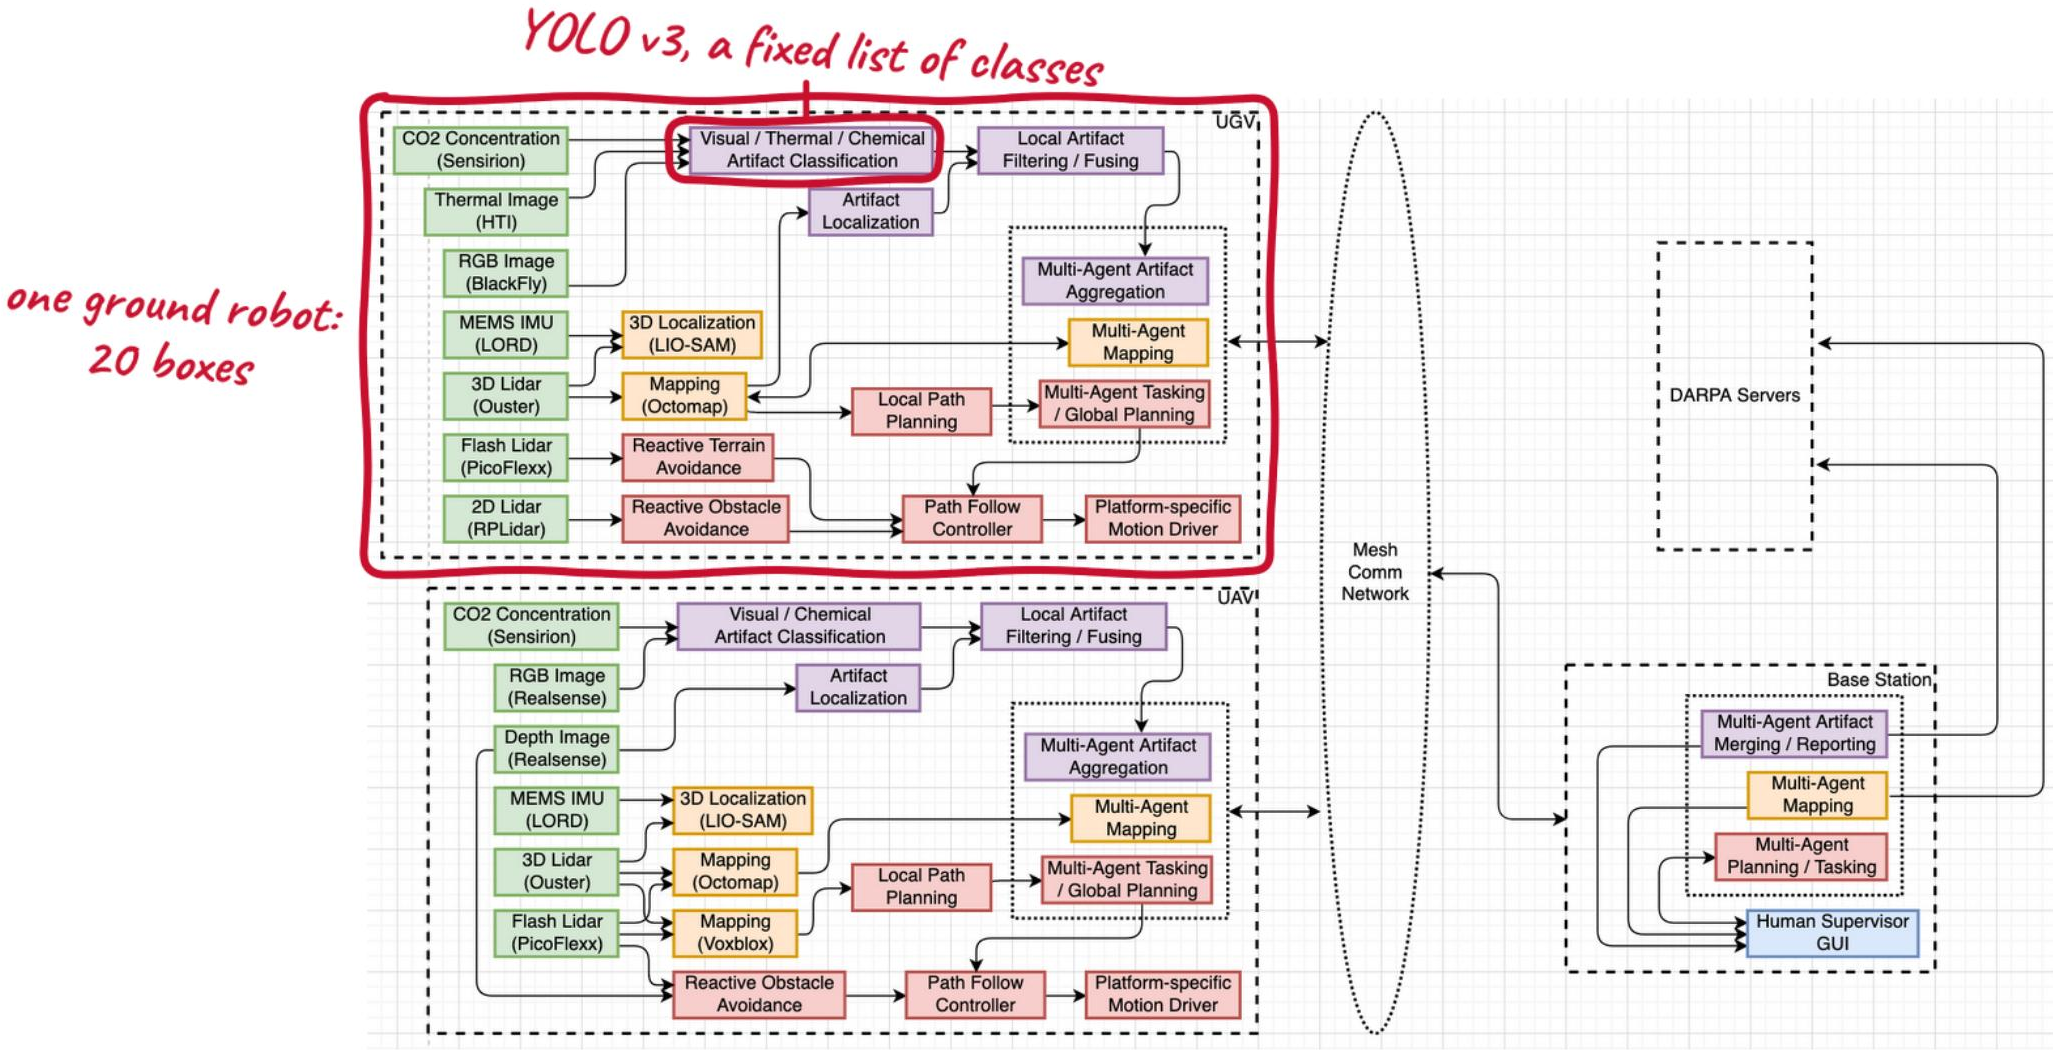
Team MARBLE

my group's robot, DARPA Subterranean Challenge, 2018 to 2021



Autonomy Stack

Team MARBLE, SubT Final debrief, 2021



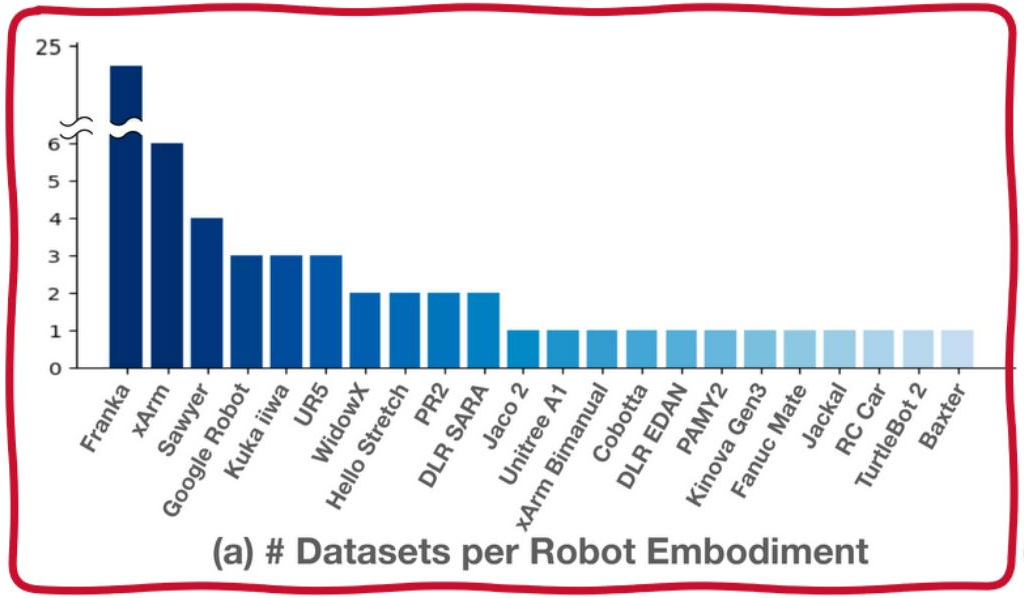
1 ACTIONS AS TOKENS

Eight Integers

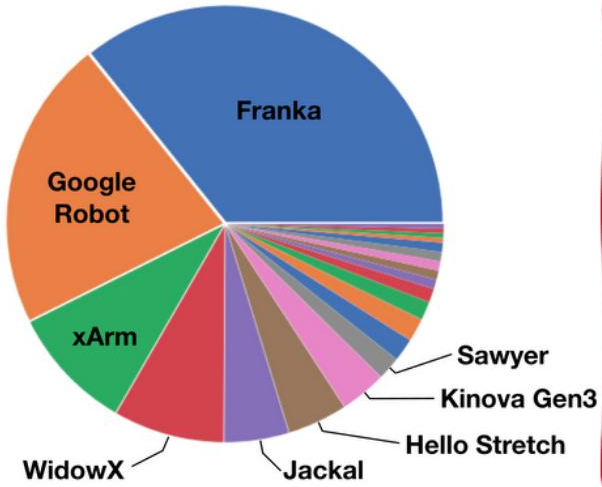
Demonstrations

Open X-Embodiment Collaboration 2023 · arXiv:2310.08864

22 kinds of robot

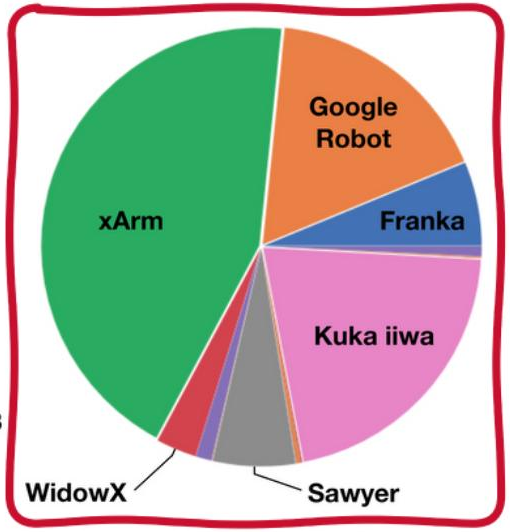


(a) # Datasets per Robot Embodiment



(b) # Scenes per Embodiment

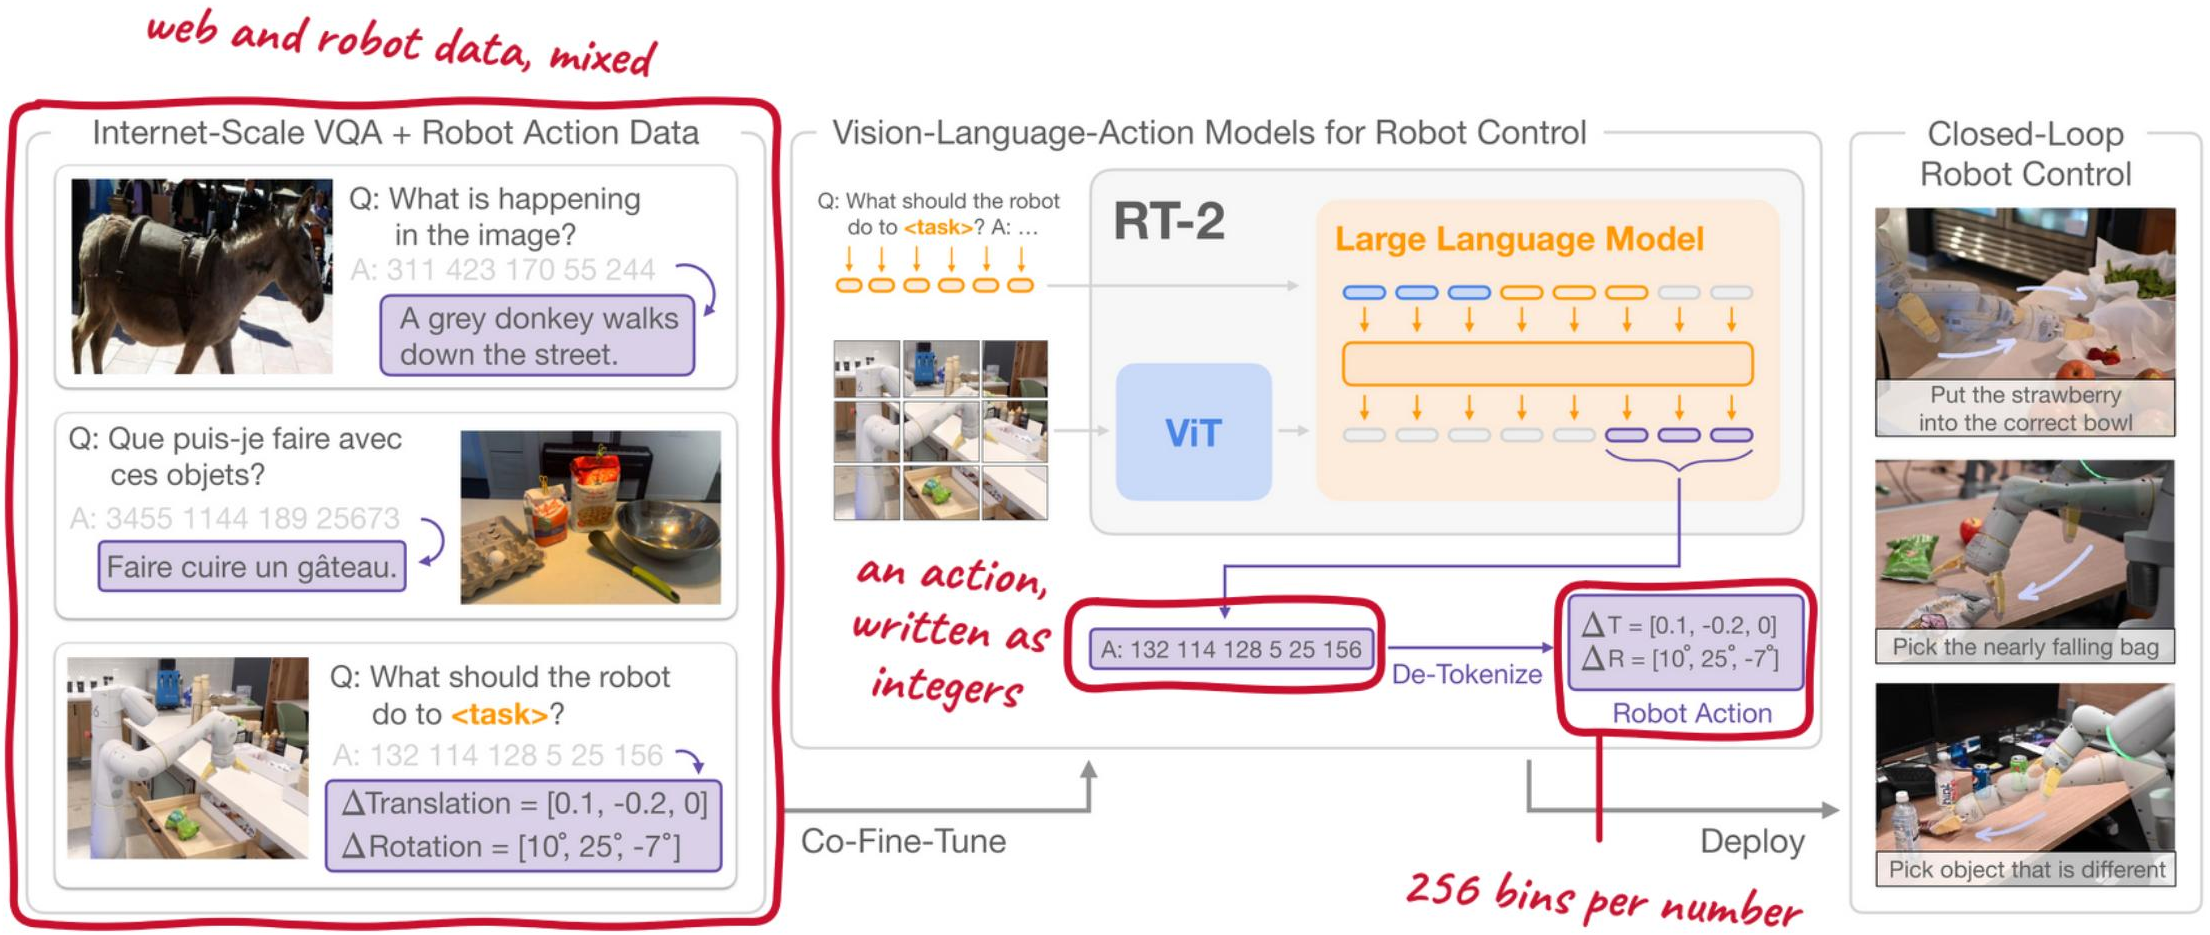
over 1M trajectories



(c) # Trajectories per Embodiment

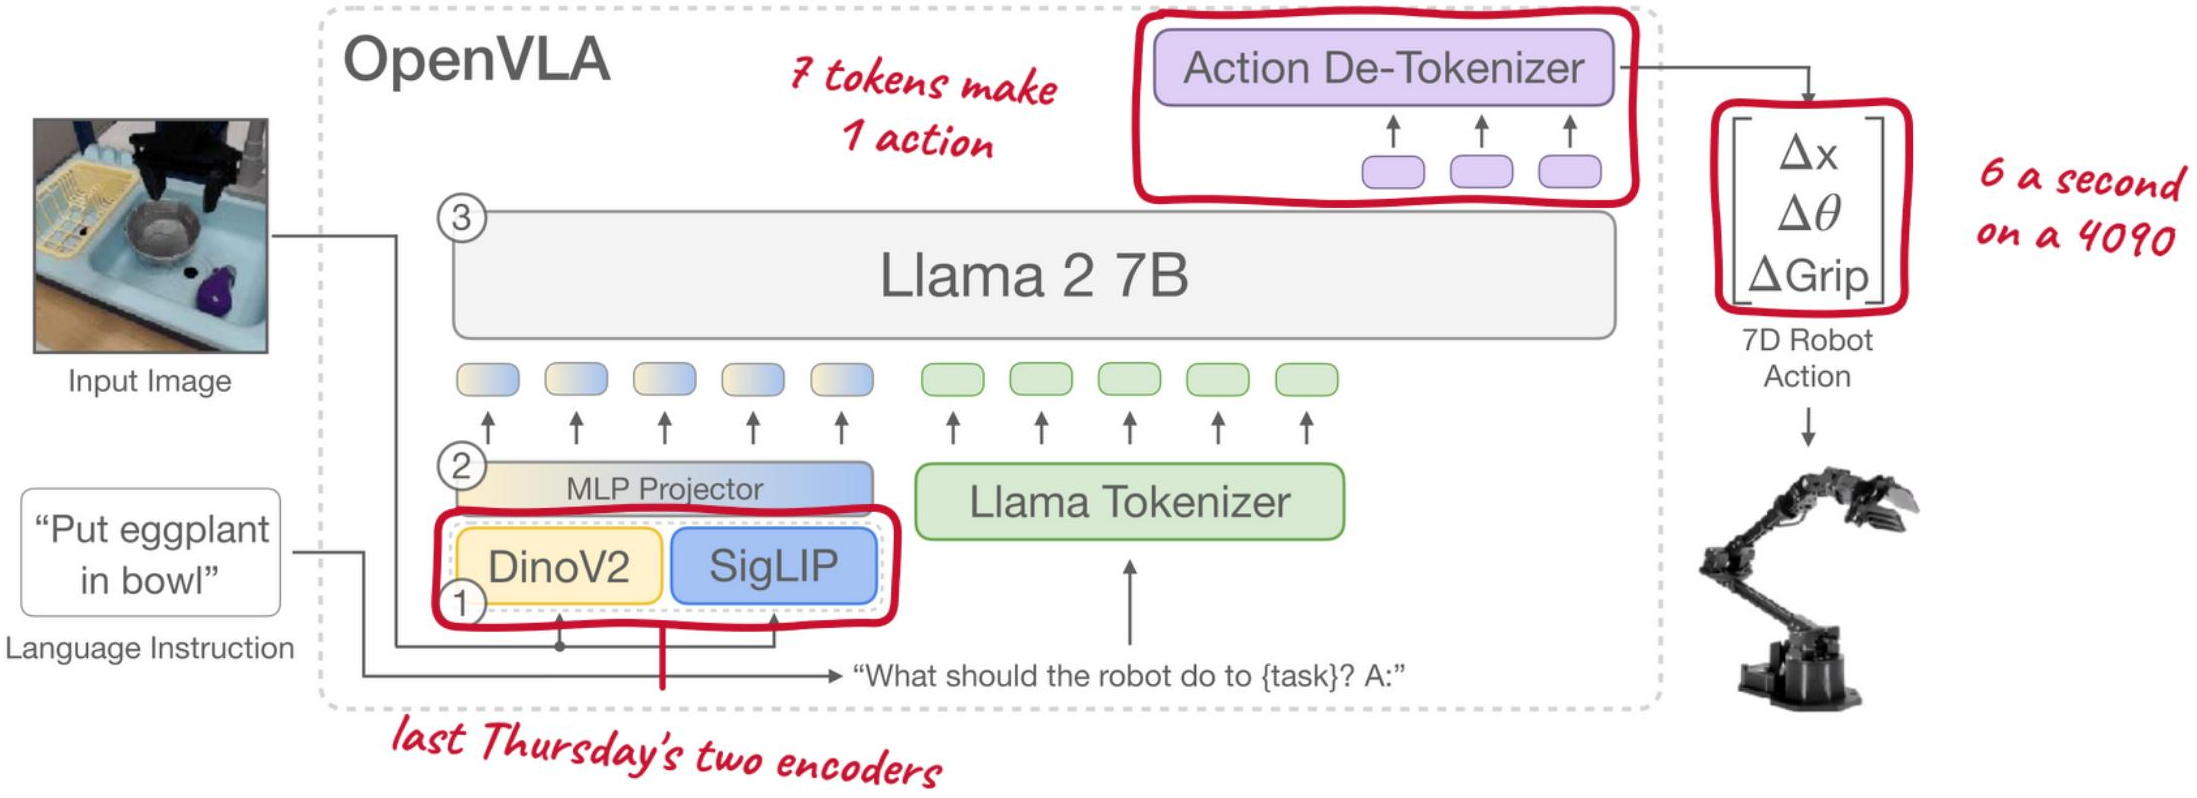
RT-2

Brohan et al. 2023 · arXiv:2307.15818



OpenVLA

Kim et al. 2024 · arXiv:2406.09246



LoRA and 4-bit

Kim et al. 2024 · arXiv:2406.09246

Strategy	Success Rate	Train Params ($\times 10^6$)	VRAM (batch 16)
Full FT	**69.7 ± 7.2 %**	7,188.1	163.3 GB*
Last layer only	30.3 ± 6.1 %	465.1	51.4 GB
Frozen vision	47.0 ± 6.9 %	6,760.4	156.2 GB*
Sandwich	62.1 ± 7.9 %	914.2	64.0 GB
LoRA, rank=32	**68.2 ± 7.5%**	**97.6**	**59.7 GB**
rank=64	**68.2 ± 7.8%**	195.2	60.5 GB

1.4% of the weights

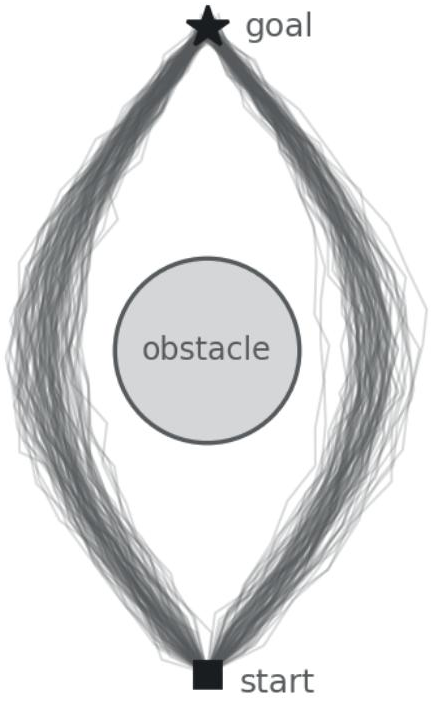
2 CONTINUOUS ACTIONS

Left or Right

Mode Averaging

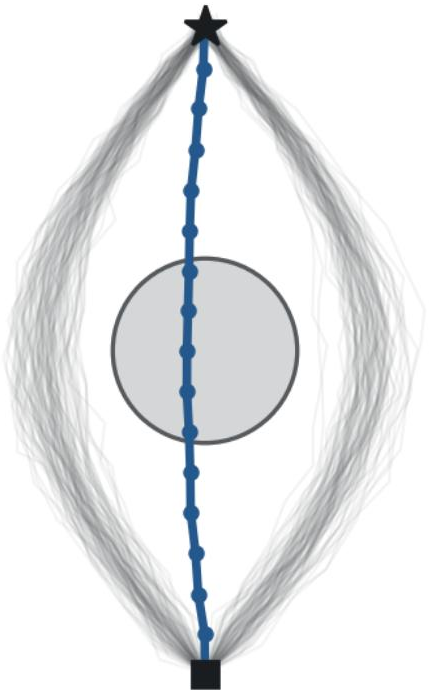
my toy run, 2,000 demonstrations

200 of the 2,000 demonstrations



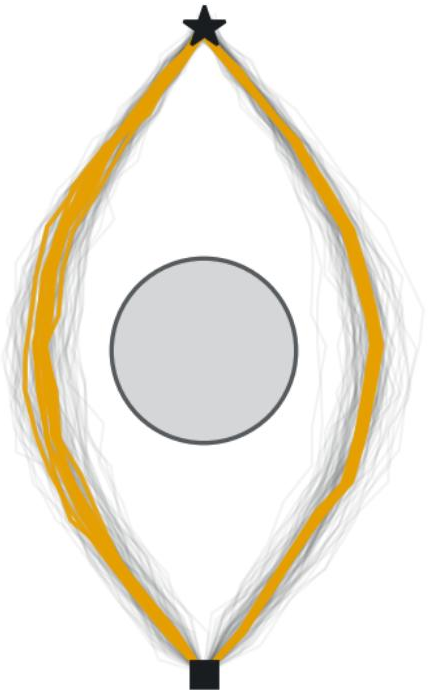
105 left, 95 right

squared-error regression



the average: through it

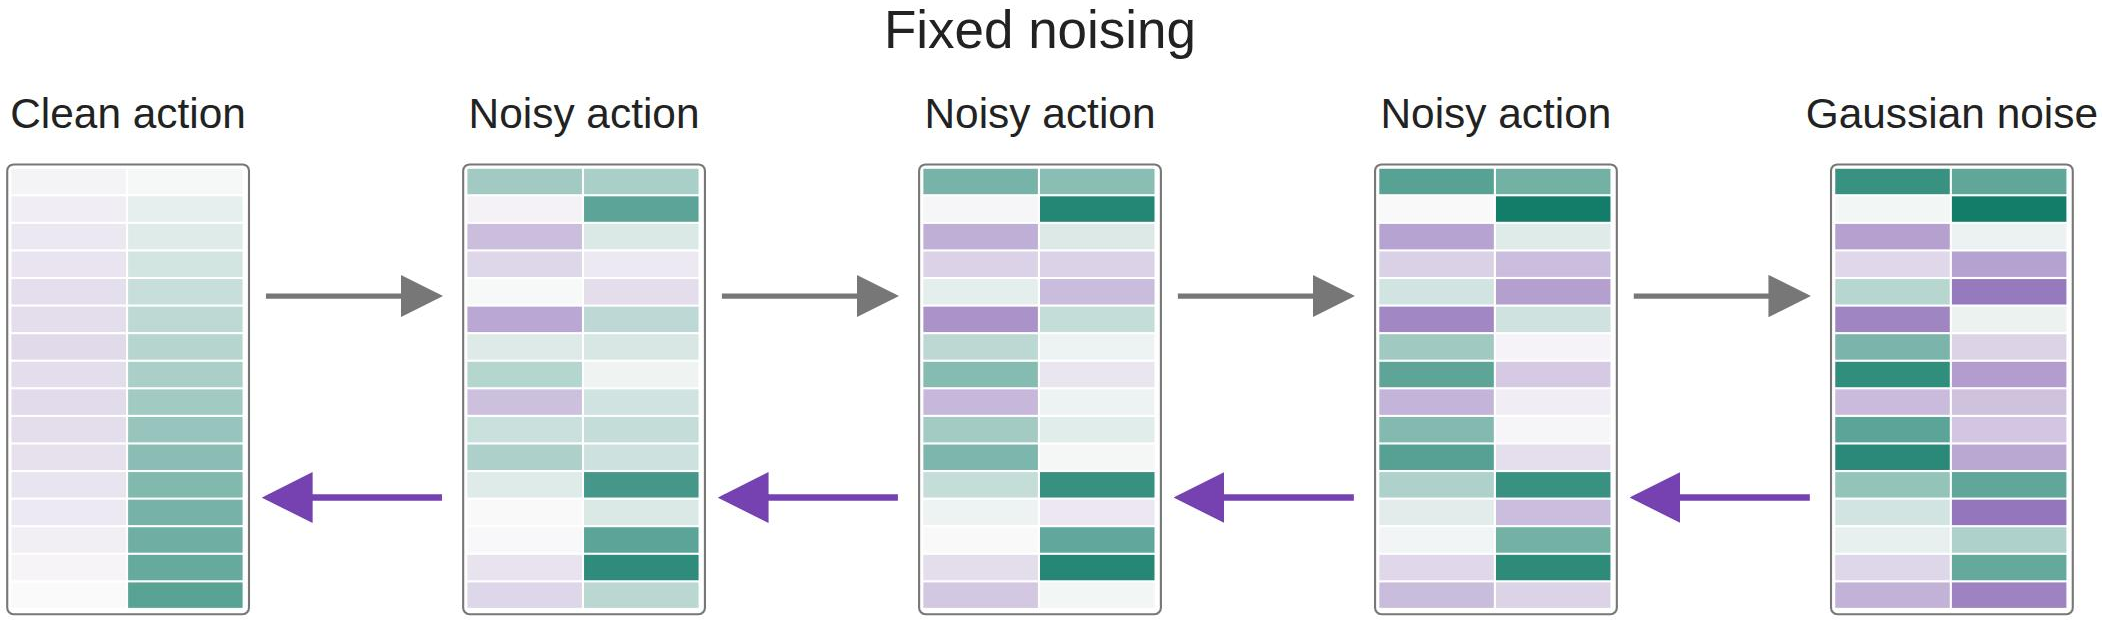
flow matching, 20 samples



10 left, 10 right

Diffusion

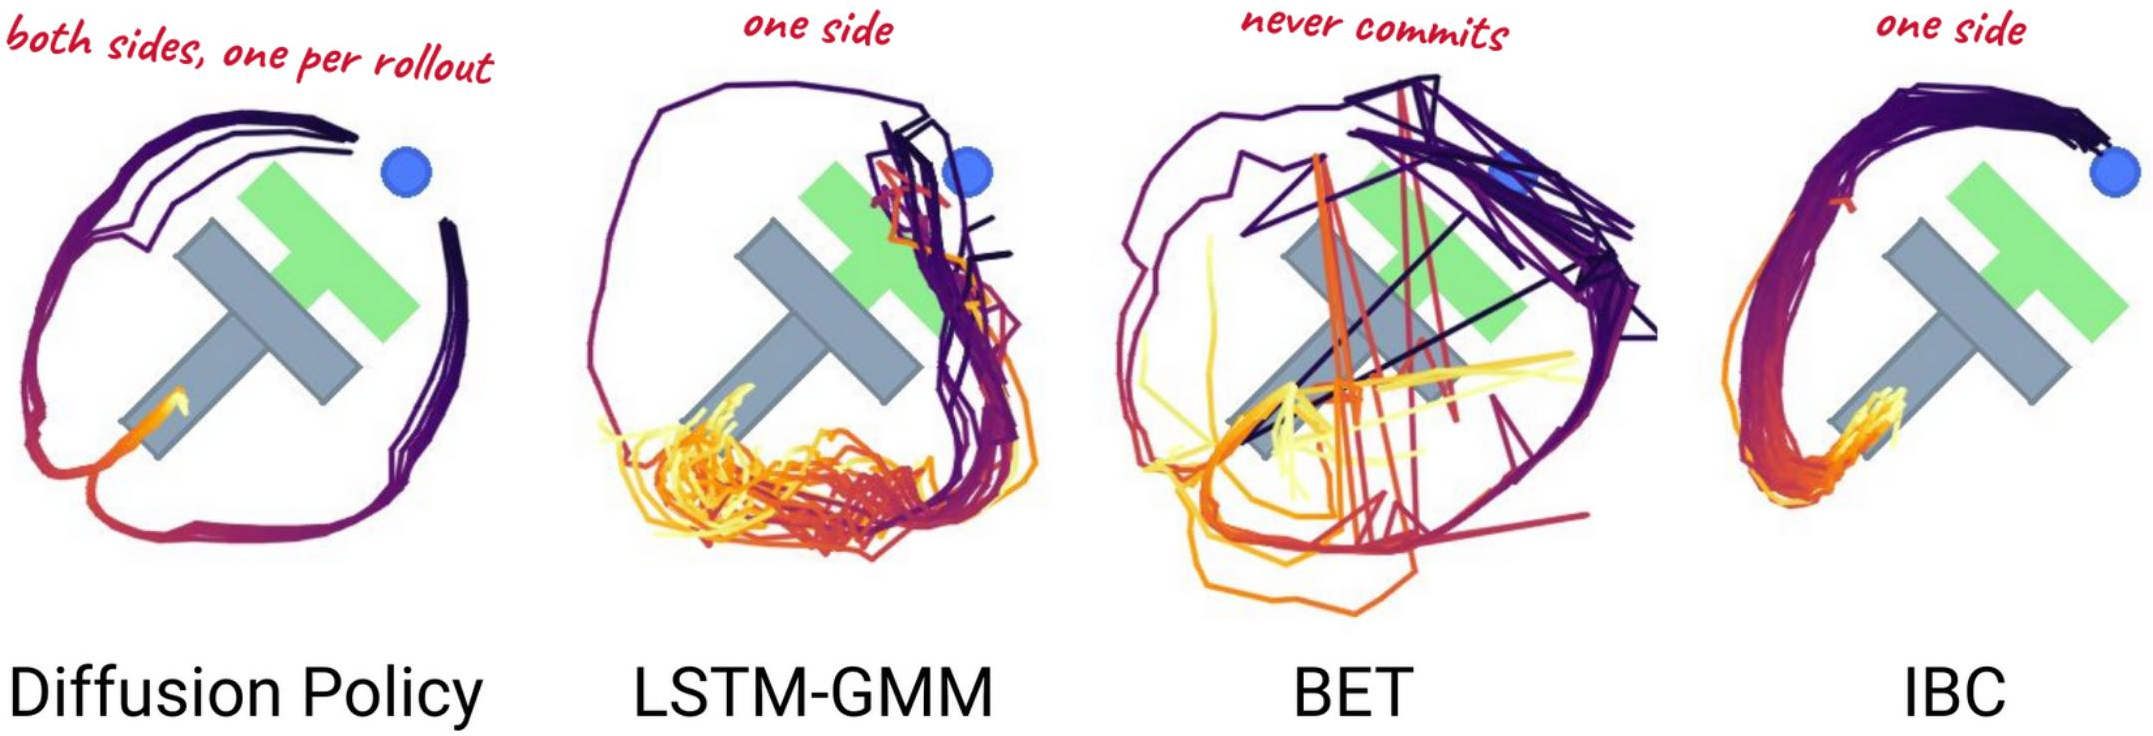
Ho et al. 2020 · arXiv:2006.11239



$$\mathcal{L}_{\text{noise}} = \mathbb{E}_{s,a,\epsilon}\left[\|\epsilon - \epsilon_\theta(a_s, s, o)\|^2\right]$$

Diffusion Policy

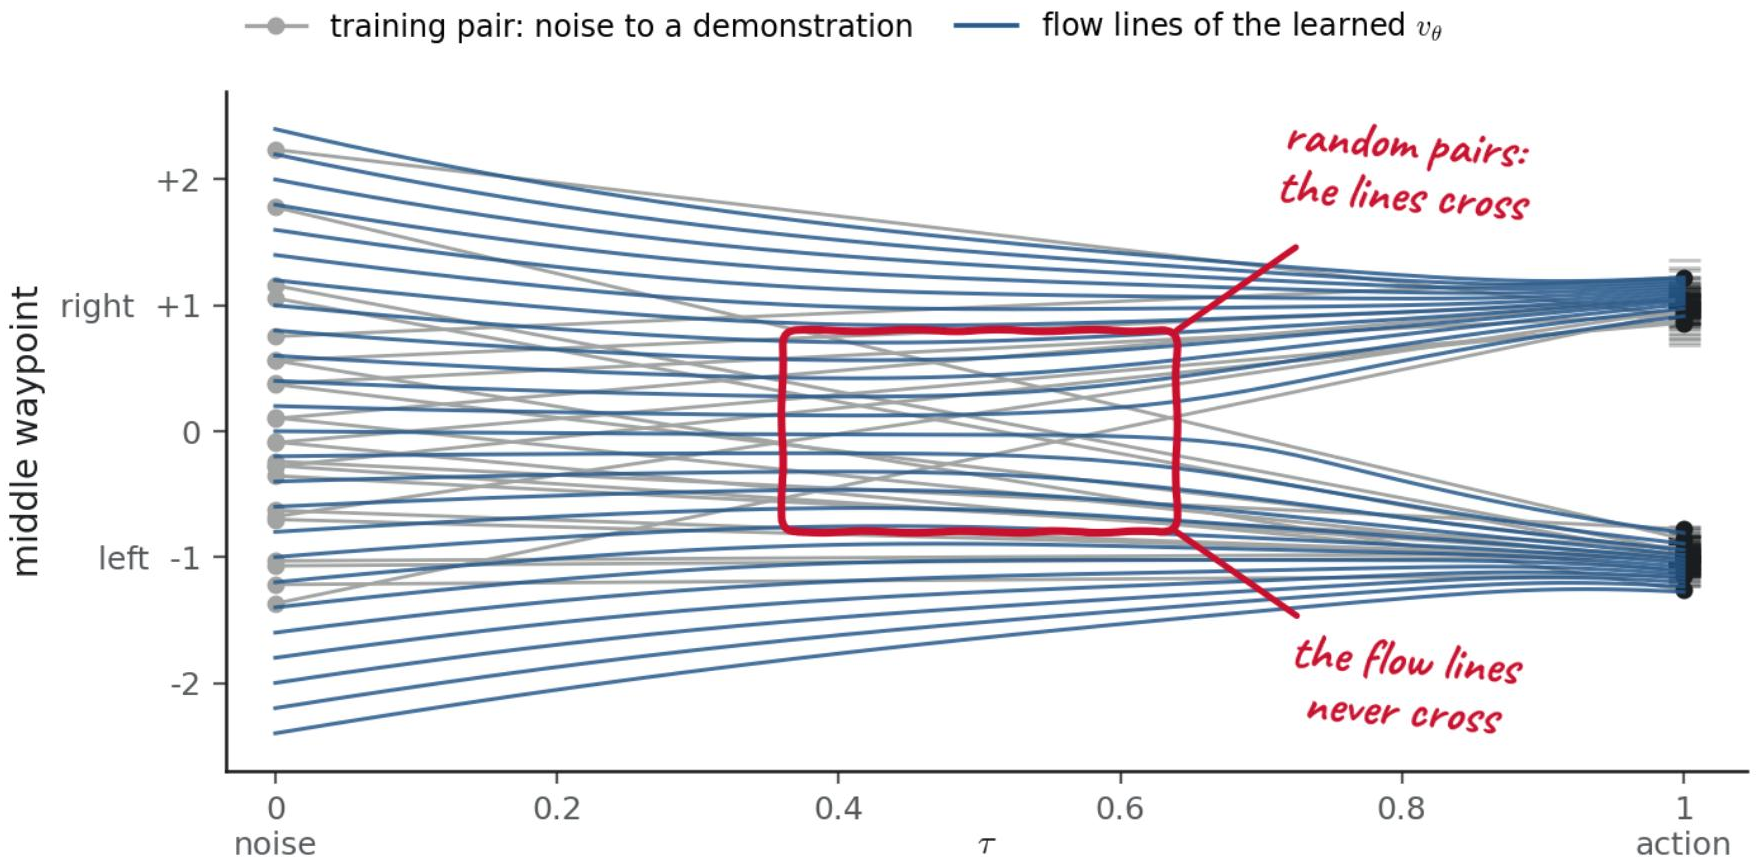
Chi et al. 2023 · arXiv:2303.04137



Flow Matching

Lipman et al. 2022 · arXiv:2210.02747

$$\frac{da^\tau}{d\tau} = v_\theta(a^\tau, \tau, o), \qquad a^0 \sim \mathcal{N}(0, I), \qquad \tau : 0 \longrightarrow 1$$

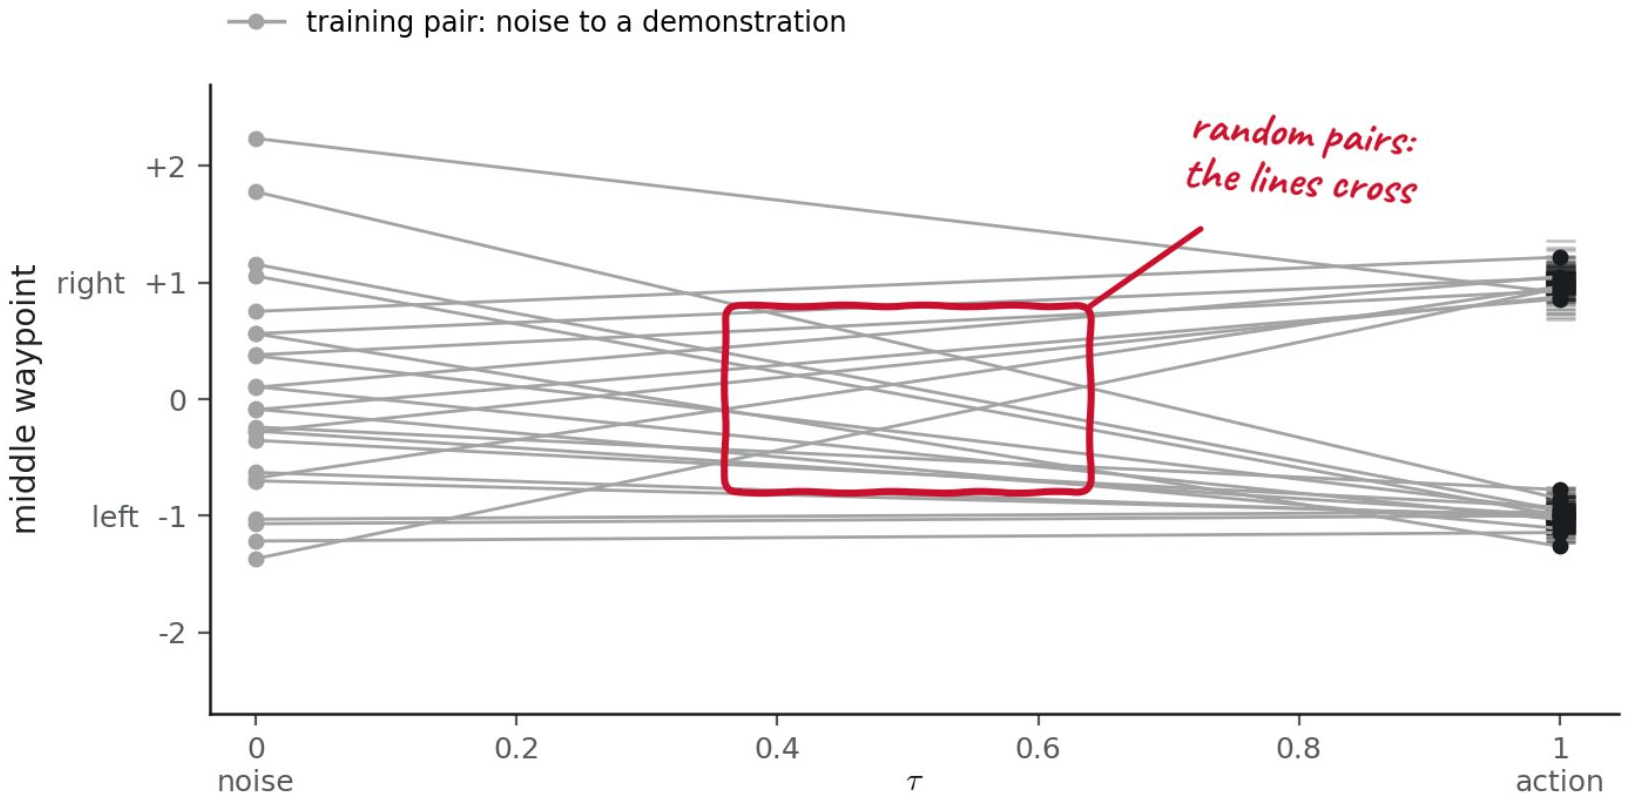


Training Pairs

Black et al. 2024 · arXiv:2410.24164

$$a^\tau = (1-\tau)\epsilon + \tau a, \qquad \frac{da^\tau}{d\tau} = a - \epsilon$$

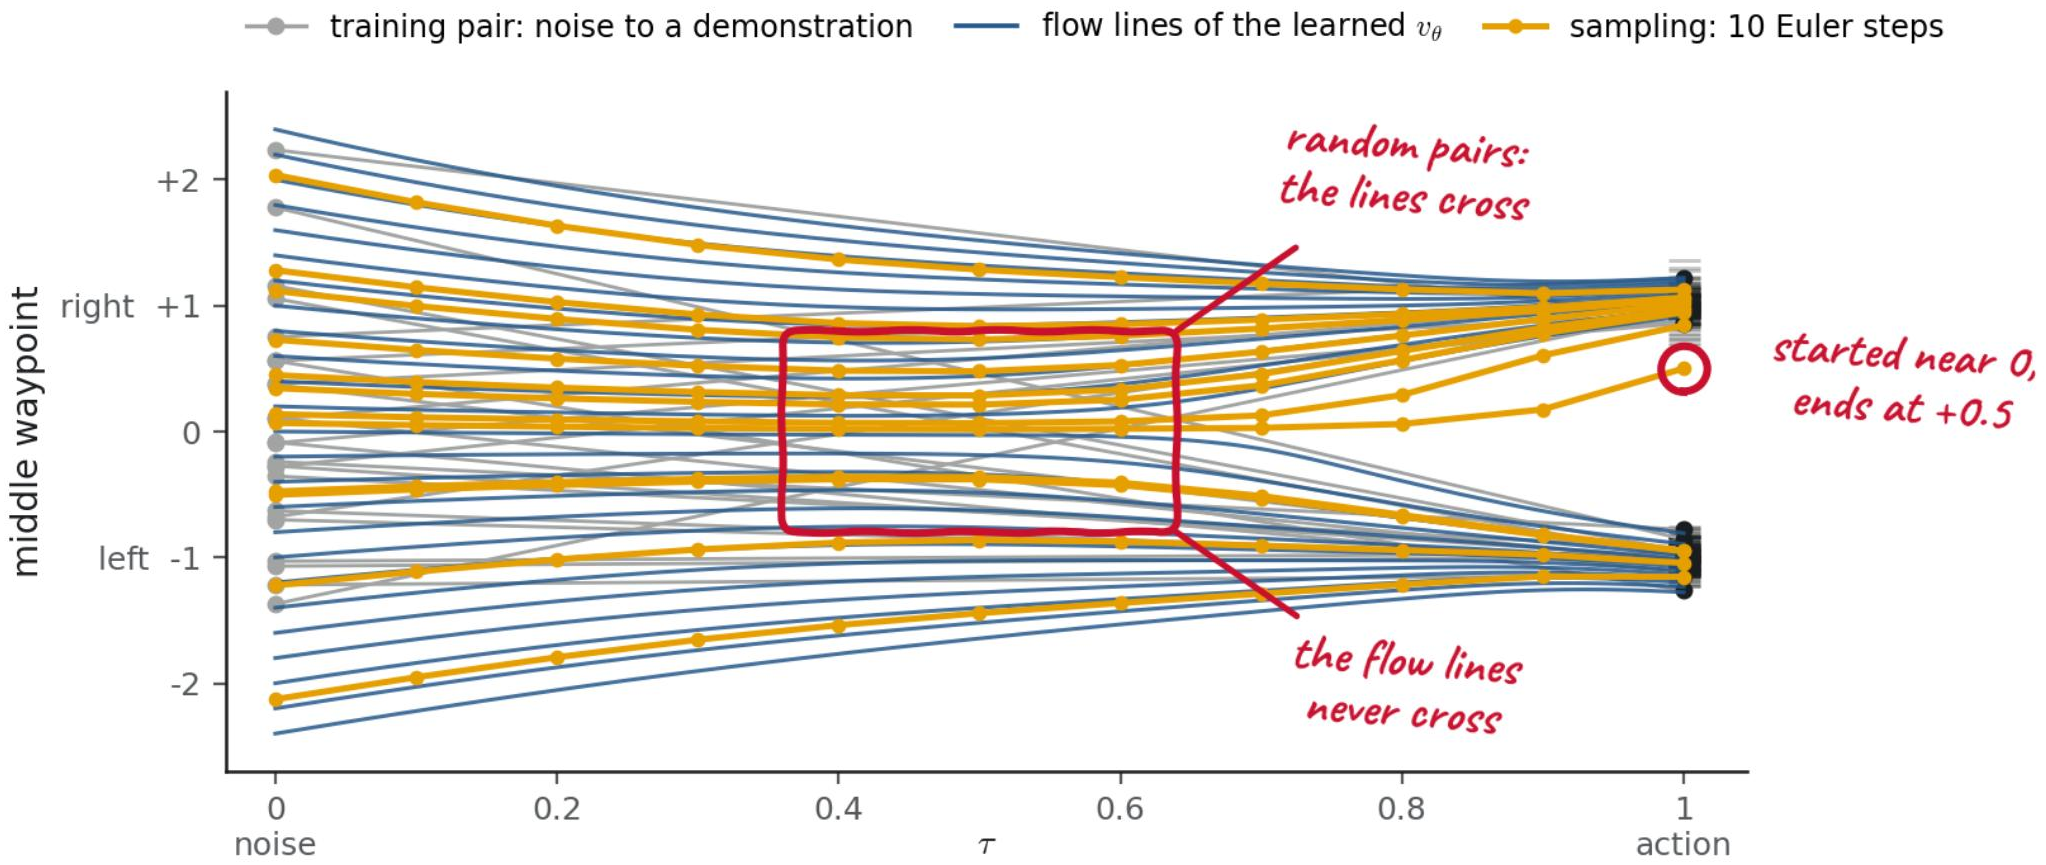
$$\mathcal{L}(\theta) = \mathbb{E}_{a,\epsilon,\tau,o}\left[\left\| v_\theta(a^\tau, \tau, o) - (a-\epsilon)\right\|^2\right]$$



Sampling

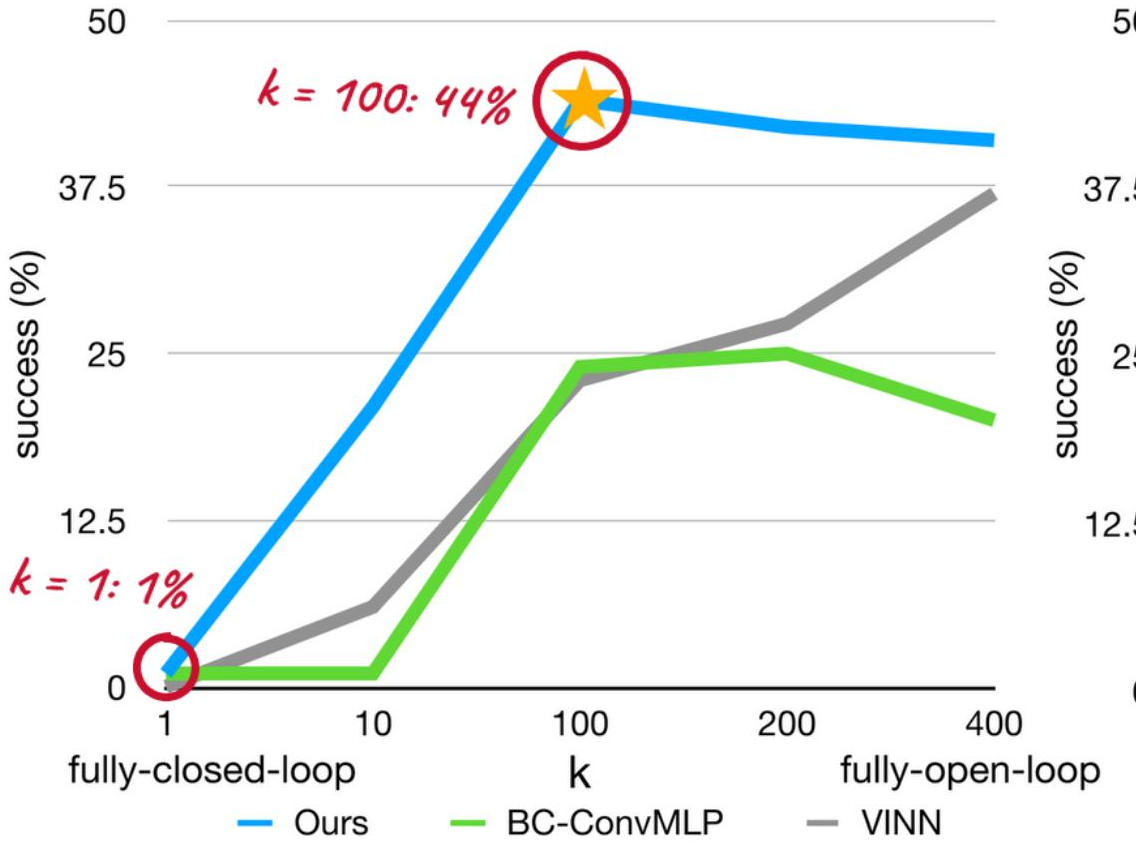
my toy run · Black et al. 2024

$$a^0 \sim \mathcal{N}(0, I), \qquad a^{\tau+\delta} = a^\tau + \delta\, v_\theta(a^\tau, \tau, o), \qquad \delta = 0.1$$

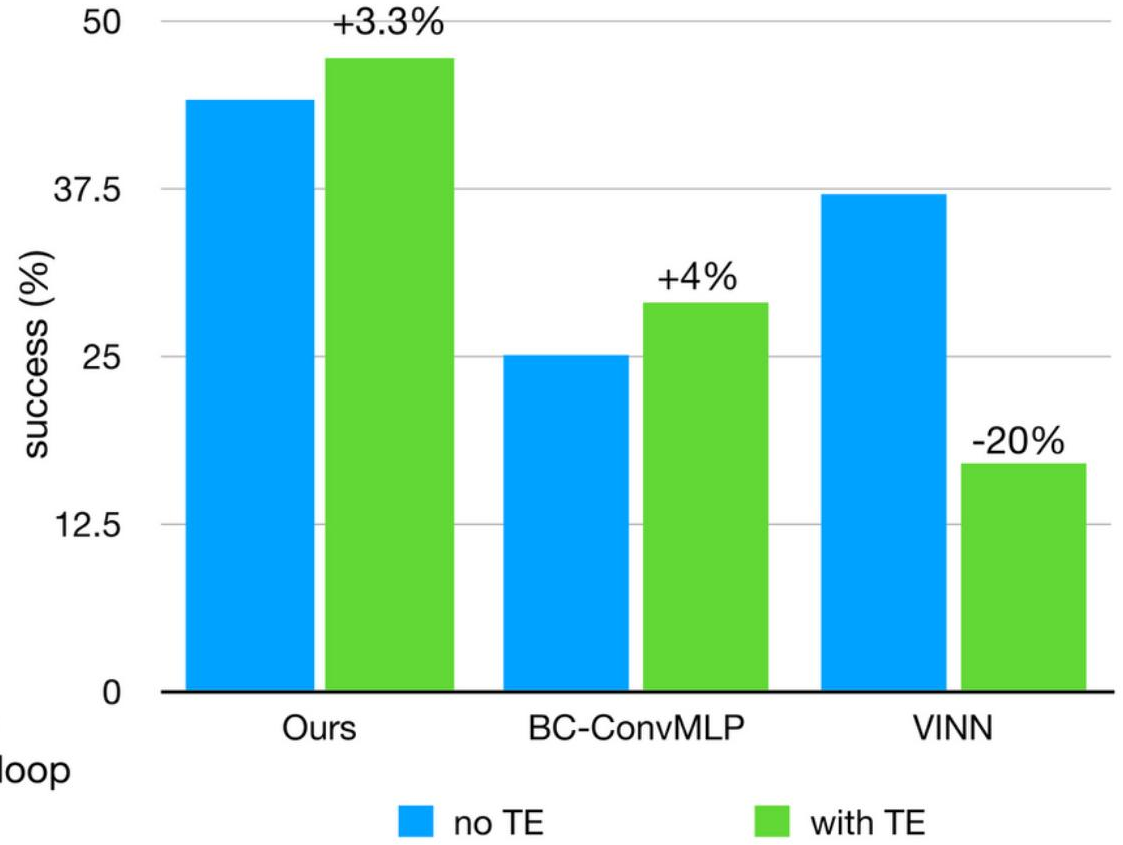


Action Chunks

Zhao et al. 2023 · arXiv:2304.13705



(a)



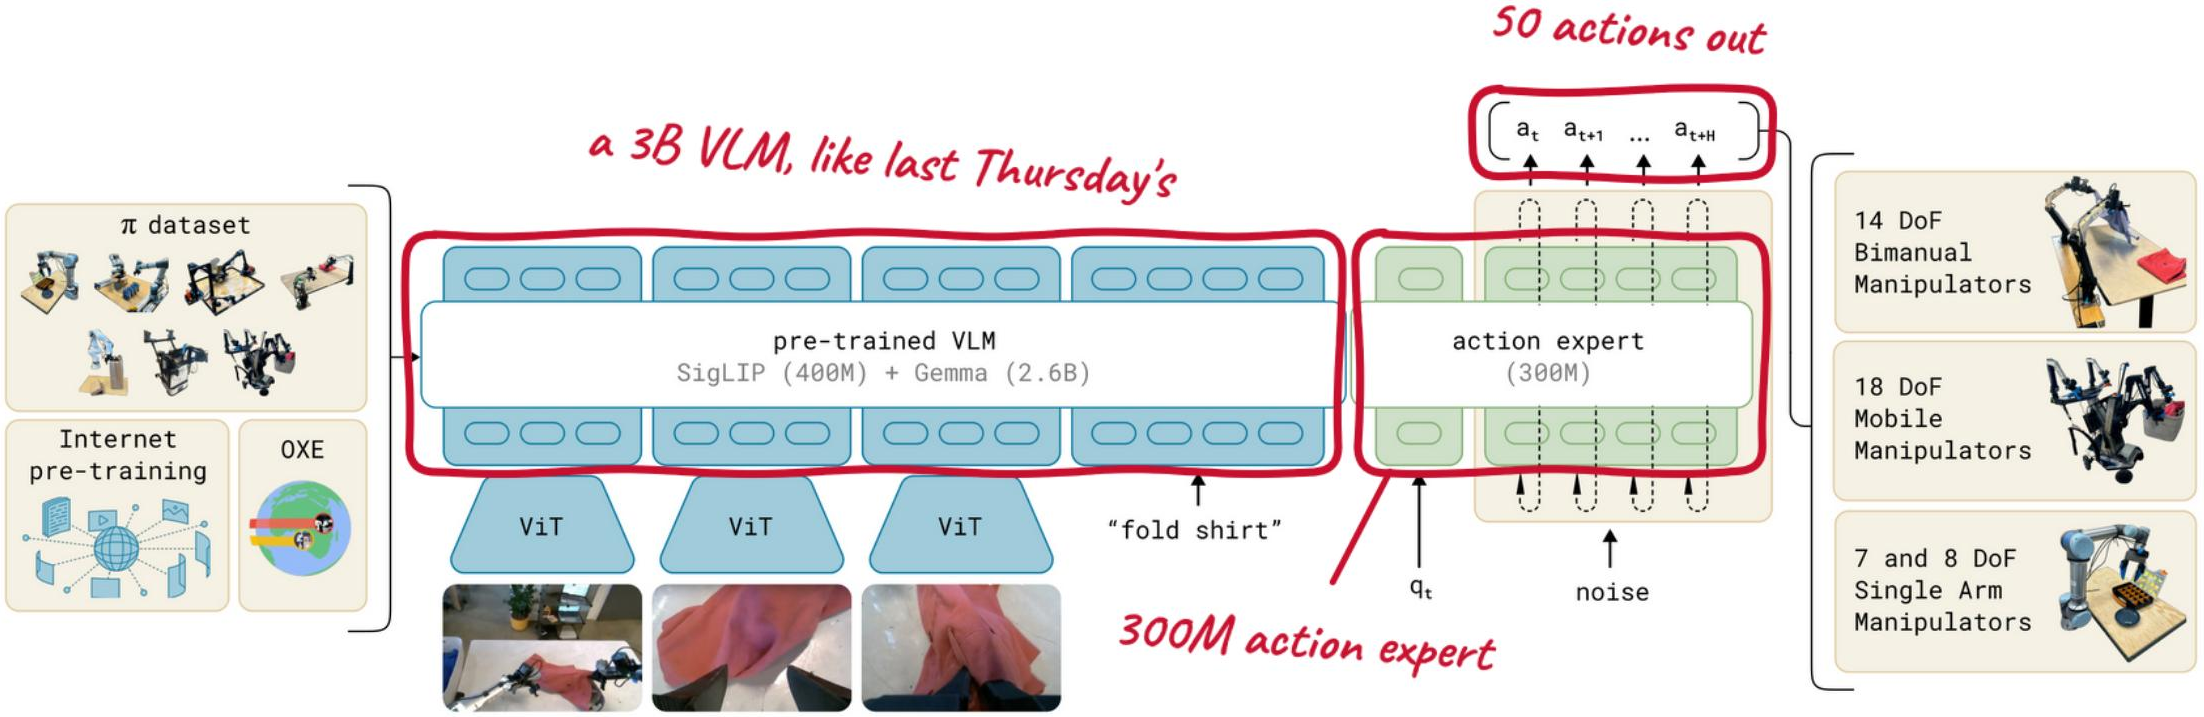
(b)

3 PHYSICAL INTELLIGENCE

An Action Expert

π_0

Black et al. 2024 · arXiv:2410.24164



73 Milliseconds

Black et al. 2024 · arXiv:2410.24164

model part	inference time
image encoders	14 ms
observation forward pass	32 ms
x10 action forward pass (flow)	27 ms
network latency (if off-board)	13 ms
total on-board inference	73 ms
total off-board inference	86 ms

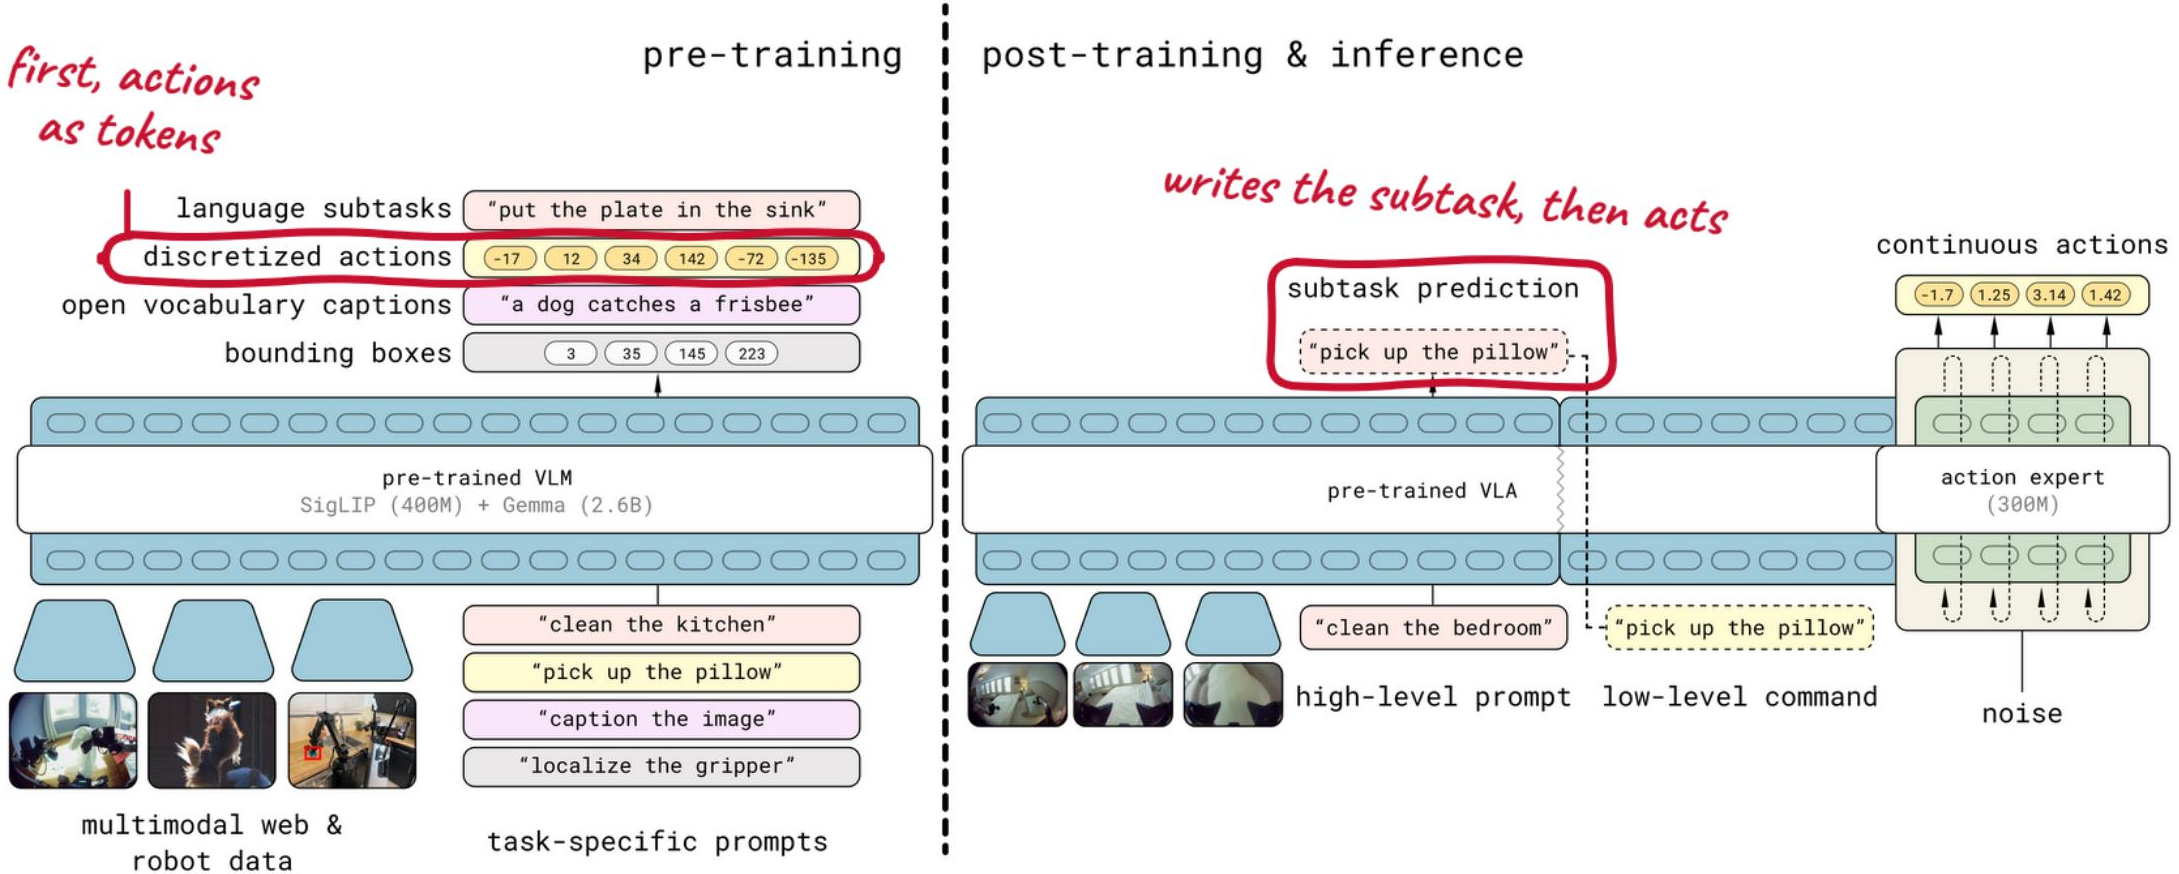
the VLM runs once: 46 ms

10 flow steps: 27 ms

50 actions every 73 ms

$\pi_{0.5}$

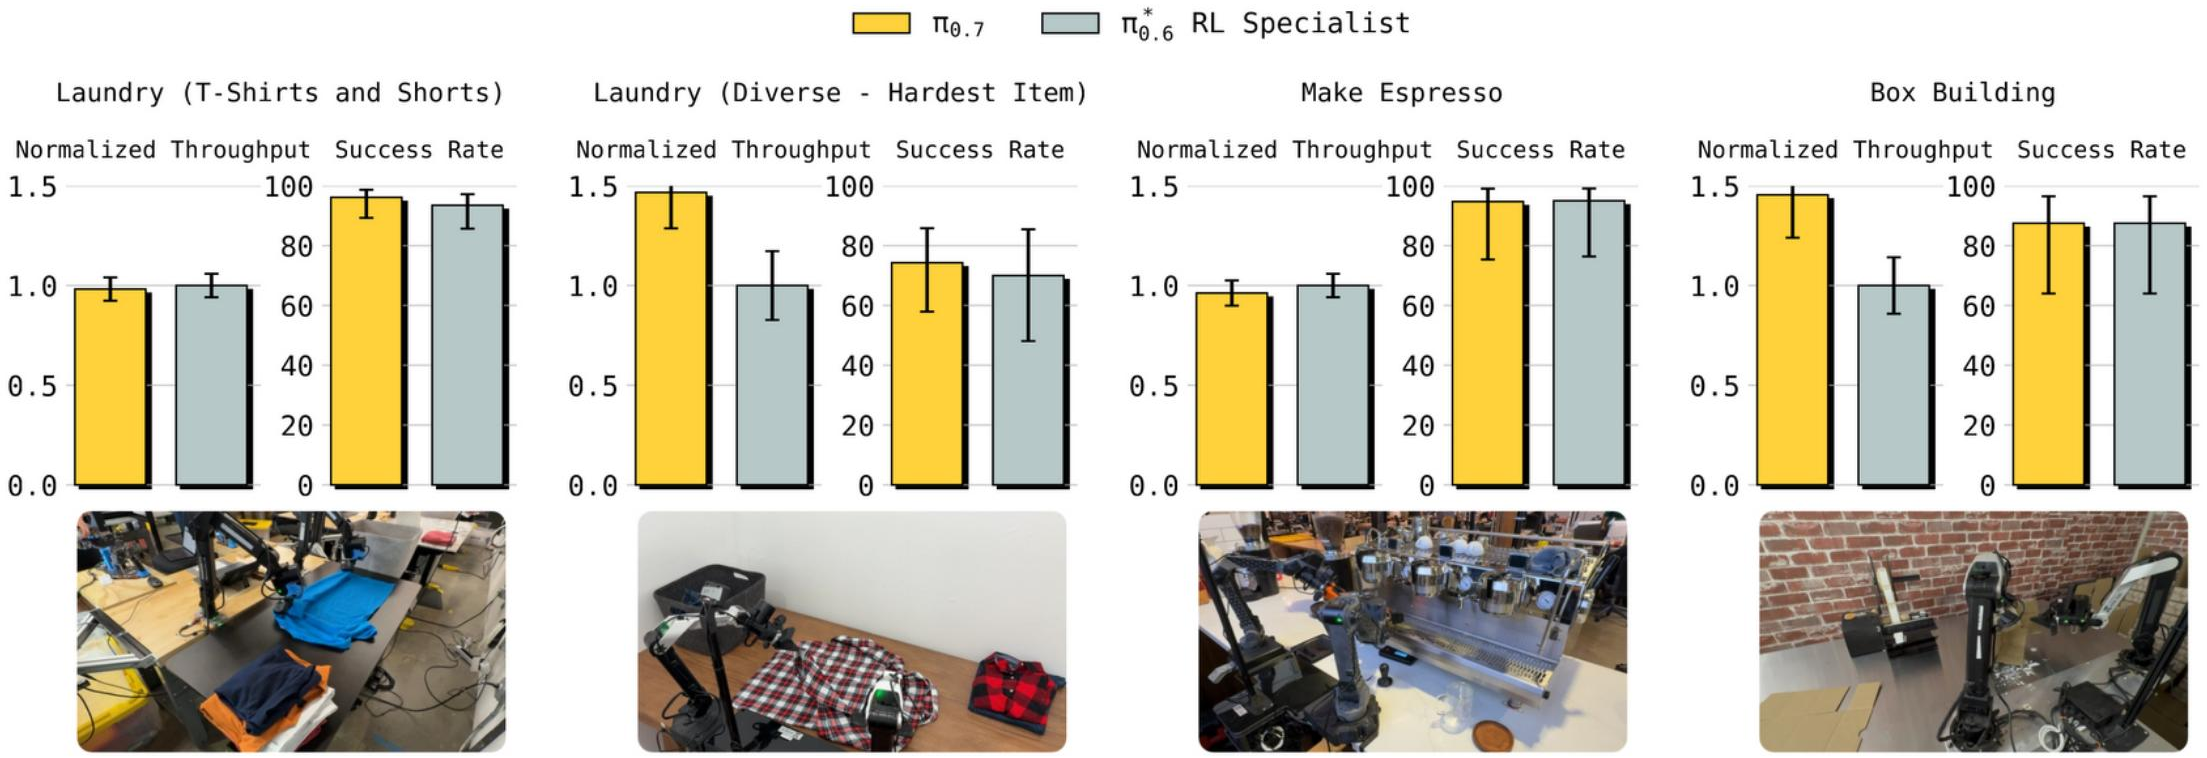
Physical Intelligence 2025 · arXiv:2504.16054



$\pi_{0.7}$

Physical Intelligence 2026 · arXiv:2604.15483

one model, level with four specialists

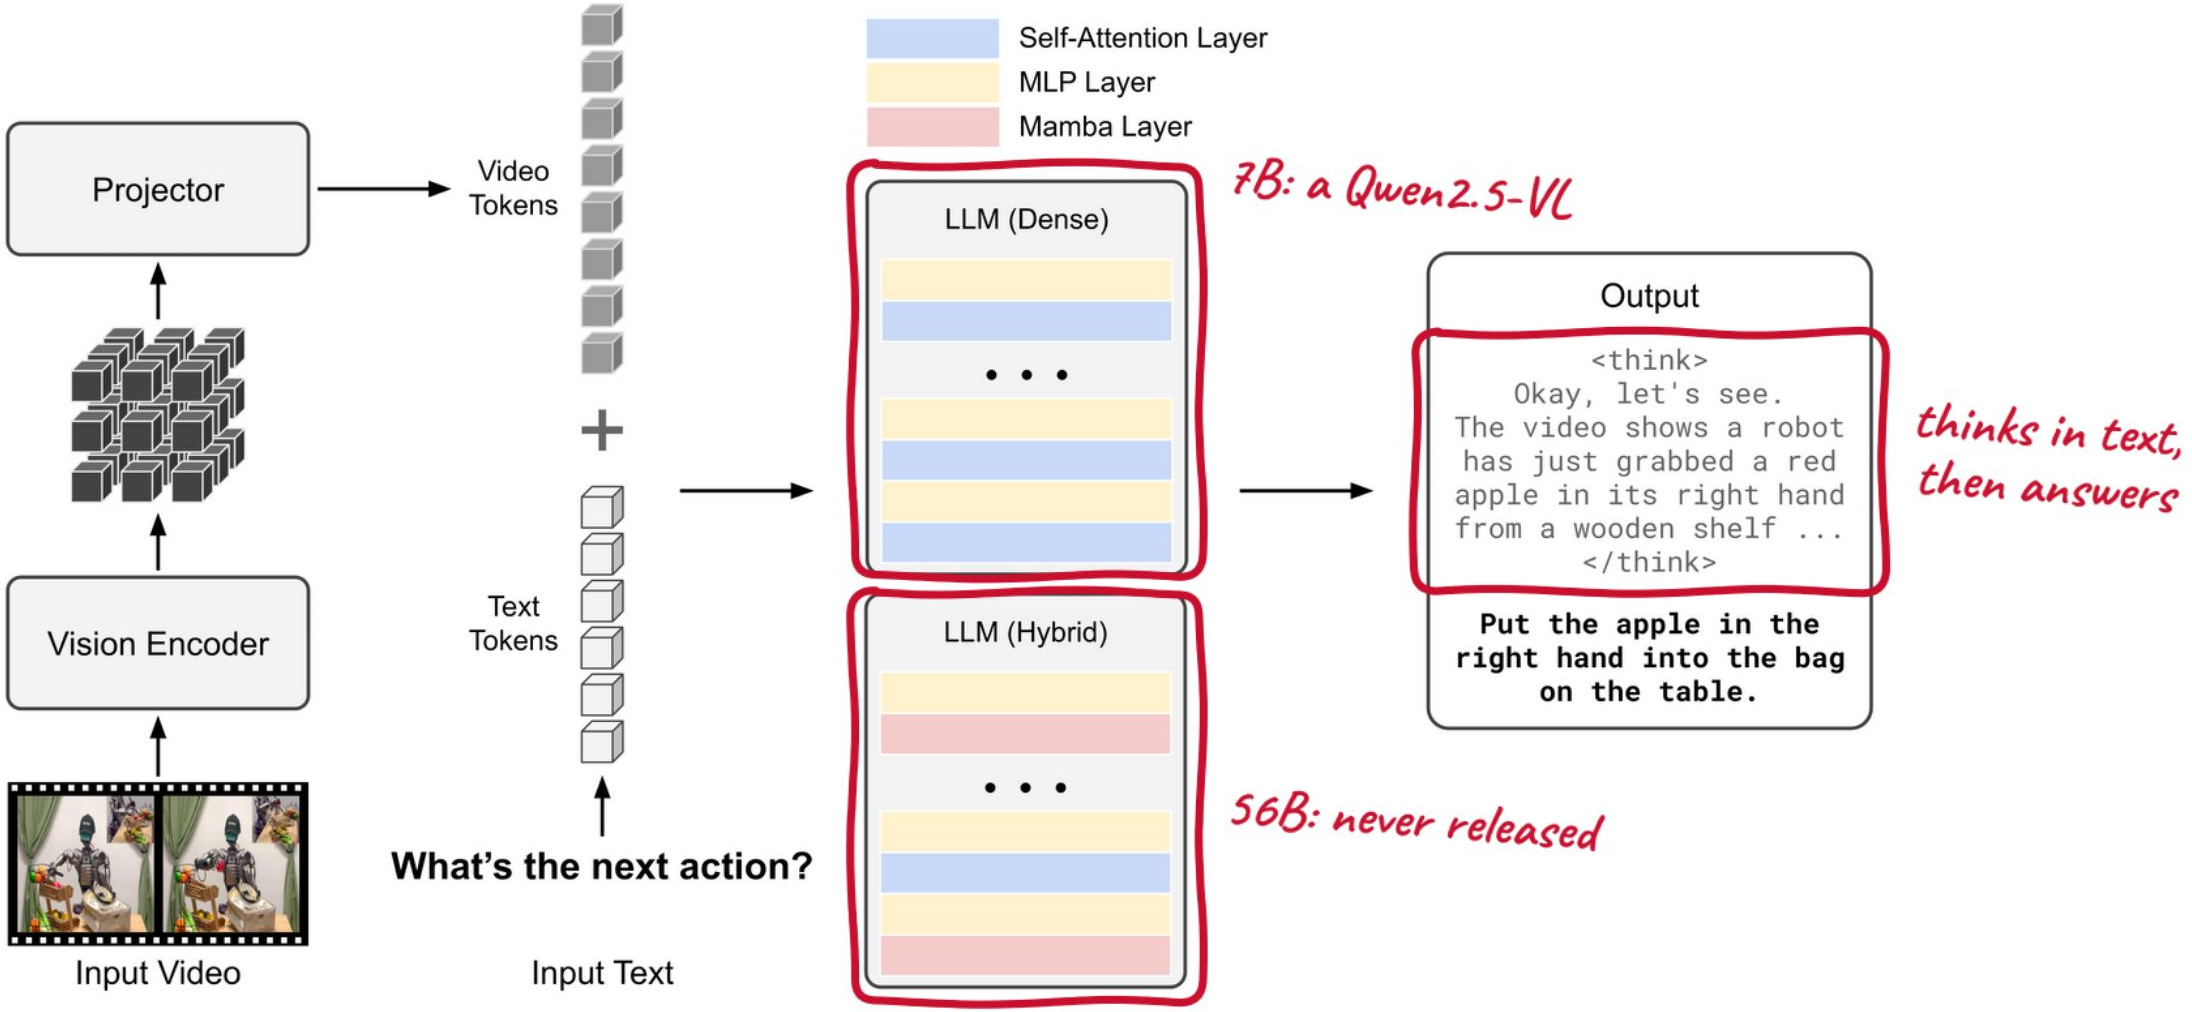


4 NVIDIA'S PHYSICAL AI STACK

A Reasoning Backbone

Cosmos Reason

NVIDIA 2025 · arXiv:2503.15558



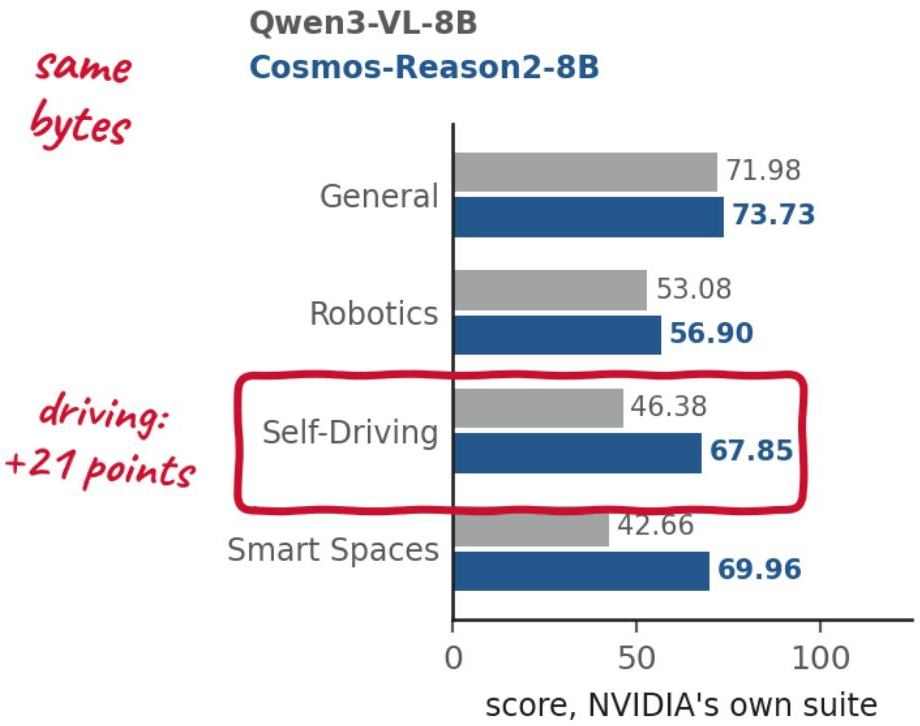
The Same Config

Cosmos-Reason2-8B and Qwen3-VL-8B, as released on Hugging Face

file	Qwen3-VL-8B	Cosmos-Reason2-8B
	git blob id, first 12 characters	
config.json	0e2df614c5fe	0e2df614c5fe
preprocessor_config.json	2ea84a437d44	2ea84a437d44
video_preprocessor_config.json	3ba673a5ad7d	3ba673a5ad7d
chat_template.json	1081bacf1af7	1081bacf1af7
tokenizer.json	c6cc1014128b	c6cc1014128b
weights, 4 shards (bytes)	17,534,339,512	17,534,339,368
parameters	8,767,123,696	8,767,123,696

same shape, new values

same bytes



driving: +21 points

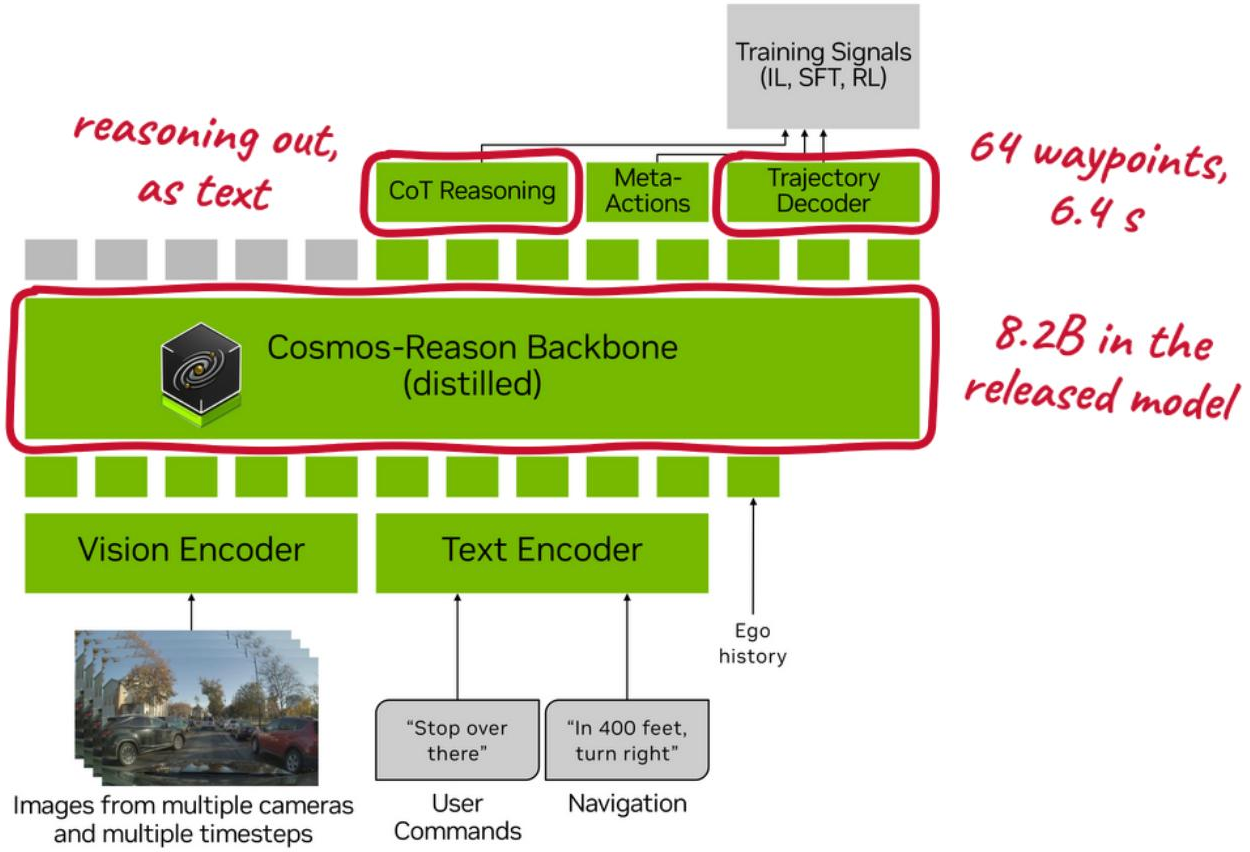
Alpamayo

Wang et al. 2025 · arXiv:2511.00088

Action: Reasoning-informed Trajectory Diffusion
- Efficient action-expert trajectory decoding via lightweight conditional flow-matching
- RL post-training to improve reasoning-action alignment and action quality

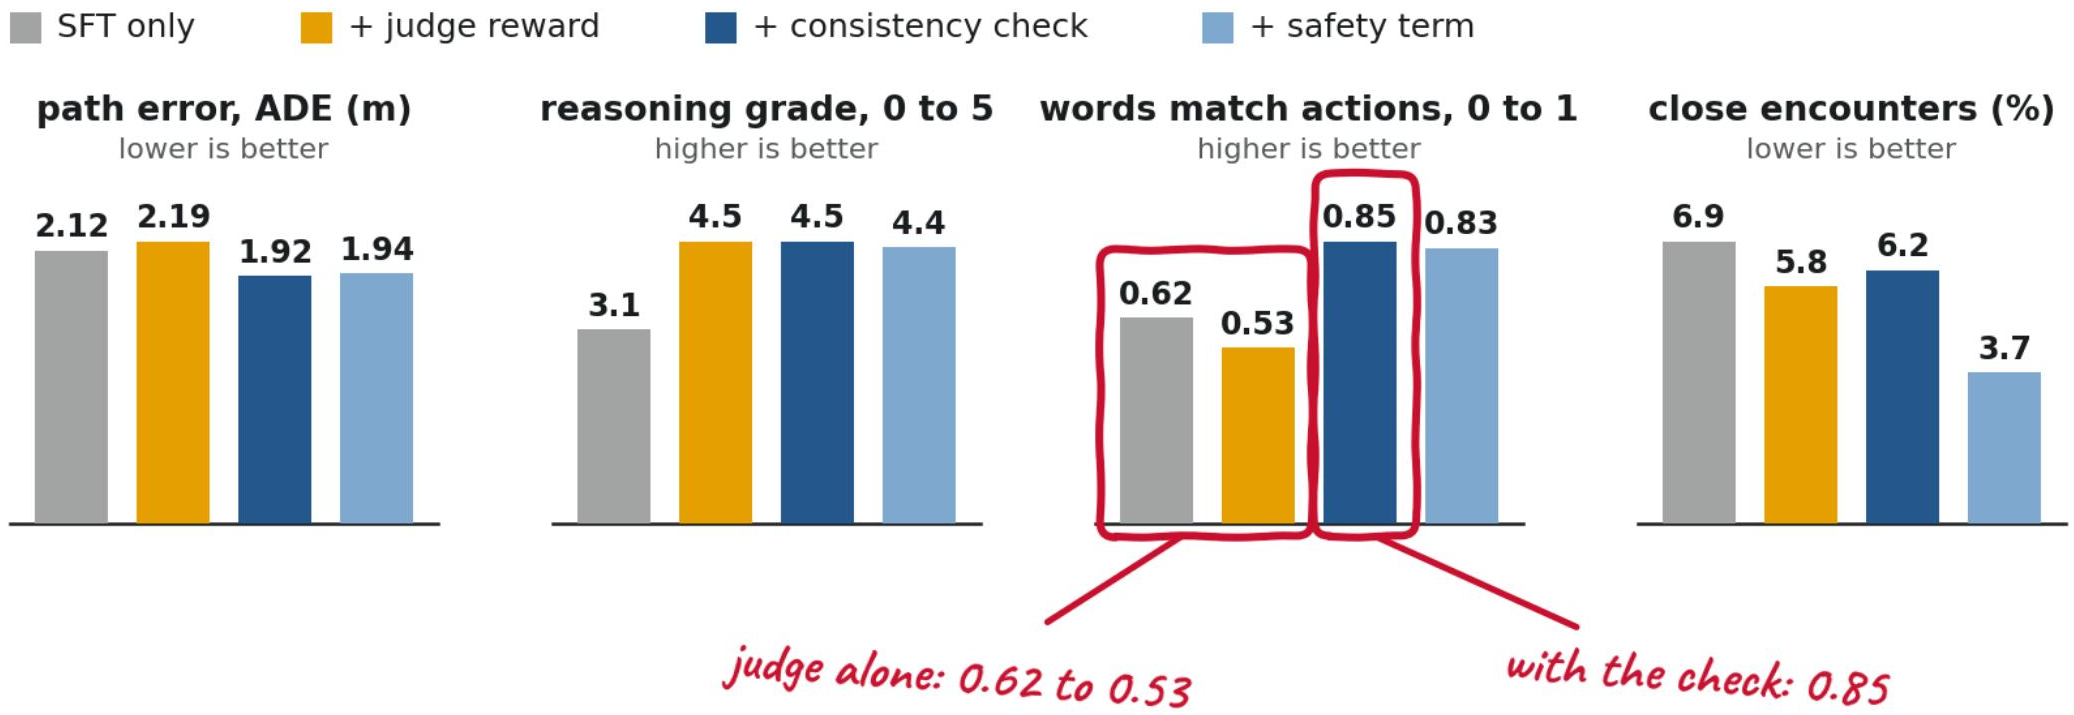
Reasoning: Internet-Pretrained Reasoning Backbone
- Pretrained to reason over internet-scale data
- RL with verifiable rewards improve causal reasoning.

Vision: Efficient Context Encoder
- Handles multiple input modalities (cameras, text)
- Efficient multi-camera, multi-timestep tokenization to reduce token sequence lengths
- Native scaling to higher resolutions and sensor counts



Chain of Causation

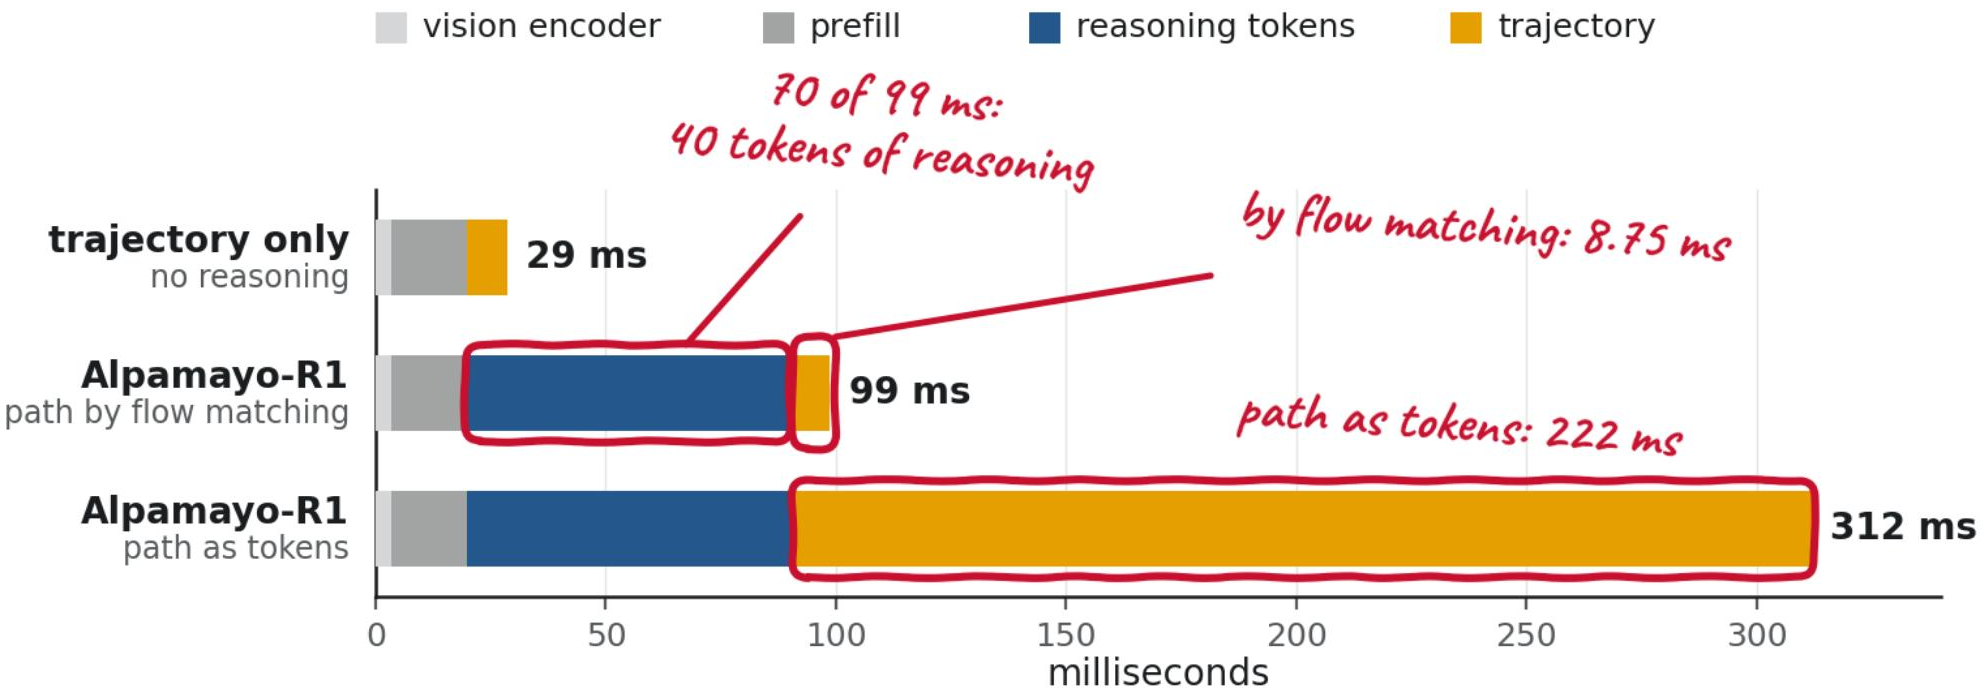
Wang et al. 2025 · arXiv:2511.00088



Training strategy	ADE ↓	Reasoning Grading ↑	Reasoning–Action Consistency Score ↑	Close Encounter Rate (%) ↓
SFT	2.12m	3.1	0.62	6.9
SFT + RL (r_{reason})	2.19m	4.5	0.53	5.8
SFT + RL ($r_{\text{reason}} + r_{\text{consistency}}$)	1.92m	4.5	0.85	6.2
SFT + RL ($r_{\text{reason}} + r_{\text{consistency}} + r_{\text{safety}}$)	1.94m	4.4	0.83	3.7

99 Milliseconds

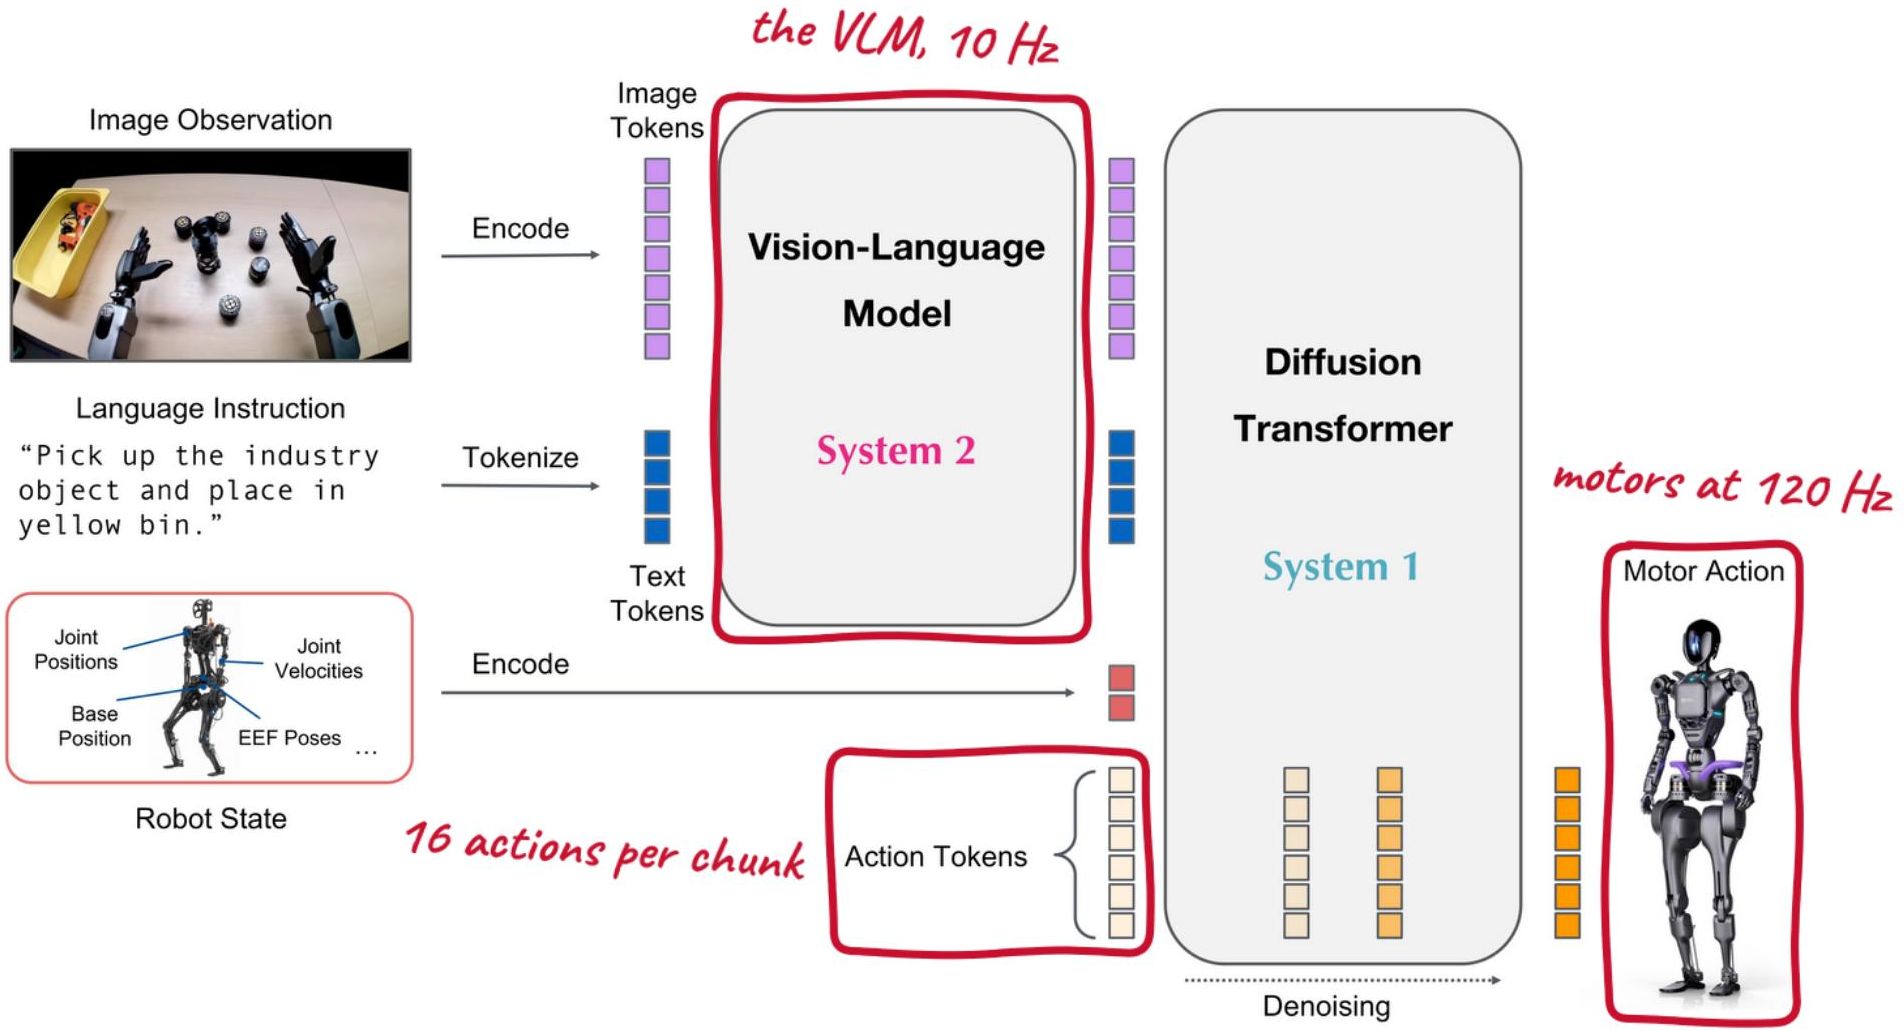
Wang et al. 2025 · arXiv:2511.00088



Model Configuration	Vision Encoder	Prefilling	Reasoning Decoding	Trajectory Decoding	Total
Baseline (trajectory-only, flow matching)	3.43ms	16.54ms	–	8.75ms (5 steps)	29ms
Alpamayo-R1 (ours, flow matching)	3.43ms	16.54ms	70ms (40 tokens)	8.75ms (5 steps)	**99ms**
Alpamayo-R1 (auto-regressive traj)	3.43ms	16.54ms	70ms (40 tokens)	222ms (127 tokens)	312ms

GR00T N1

NVIDIA 2025 · arXiv:2503.14734

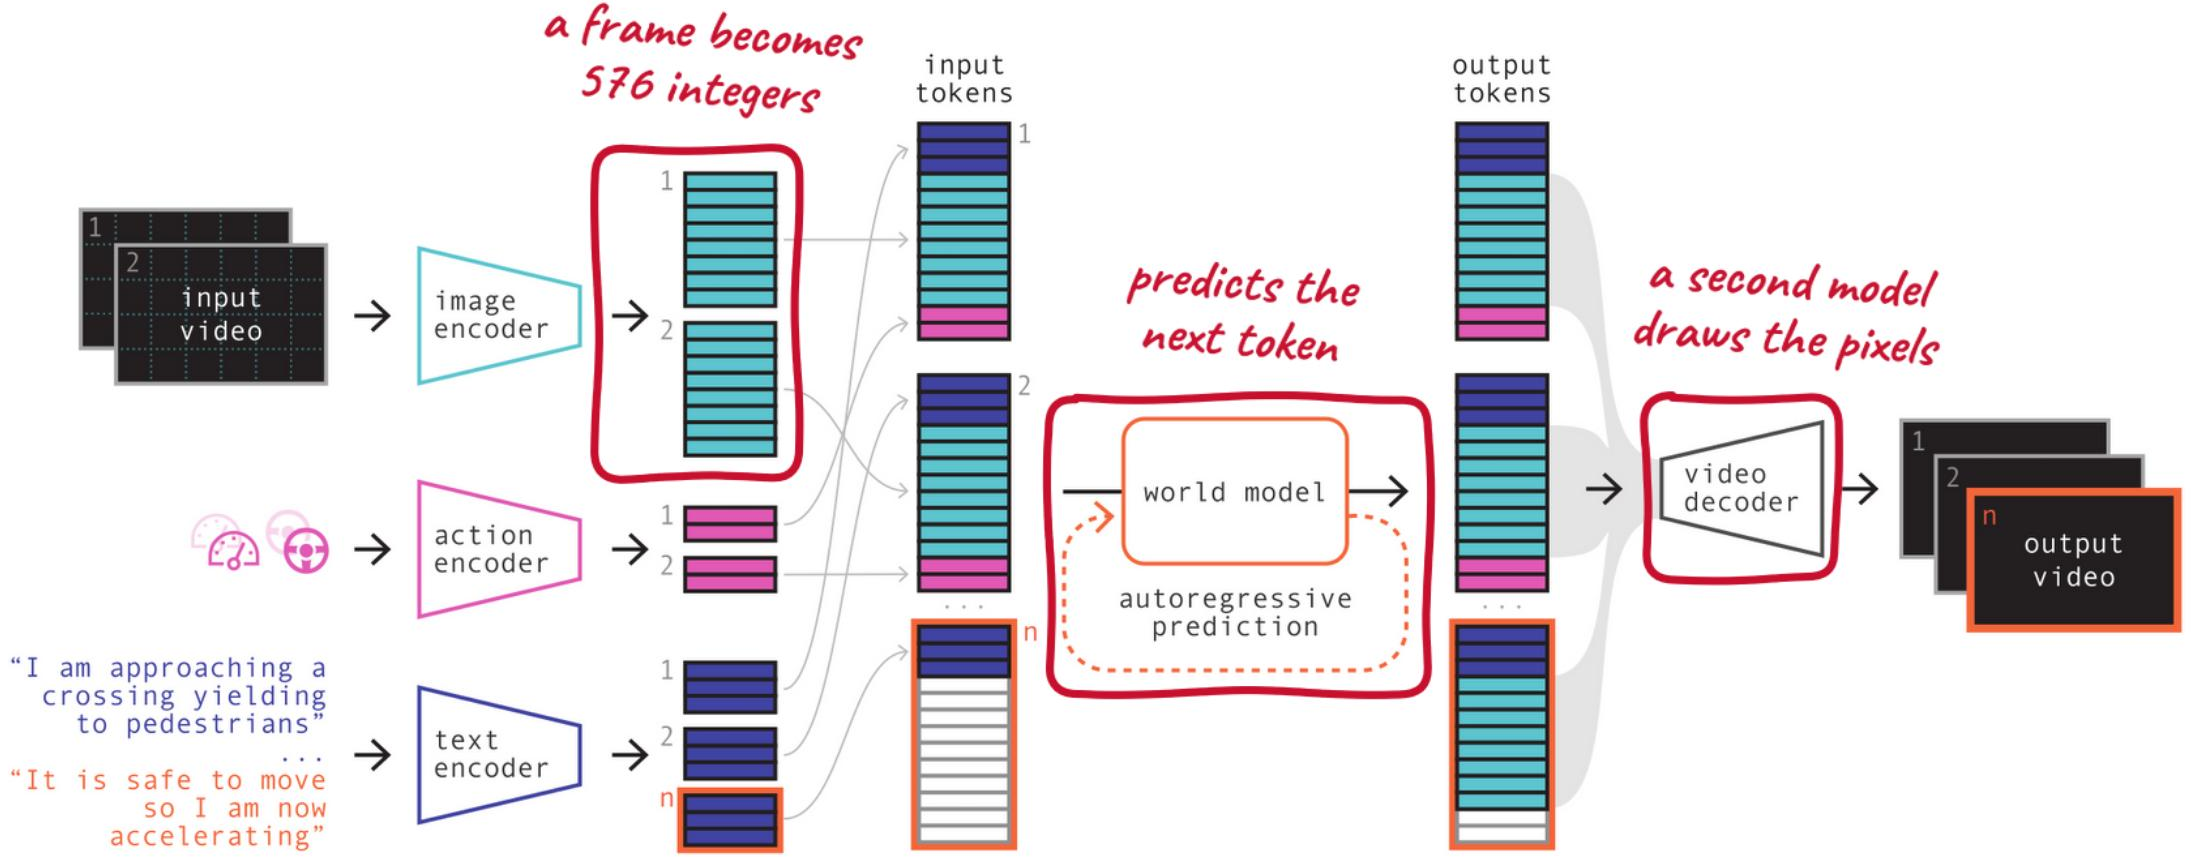


5 WORLD MODELS

Predicting the Next Frame

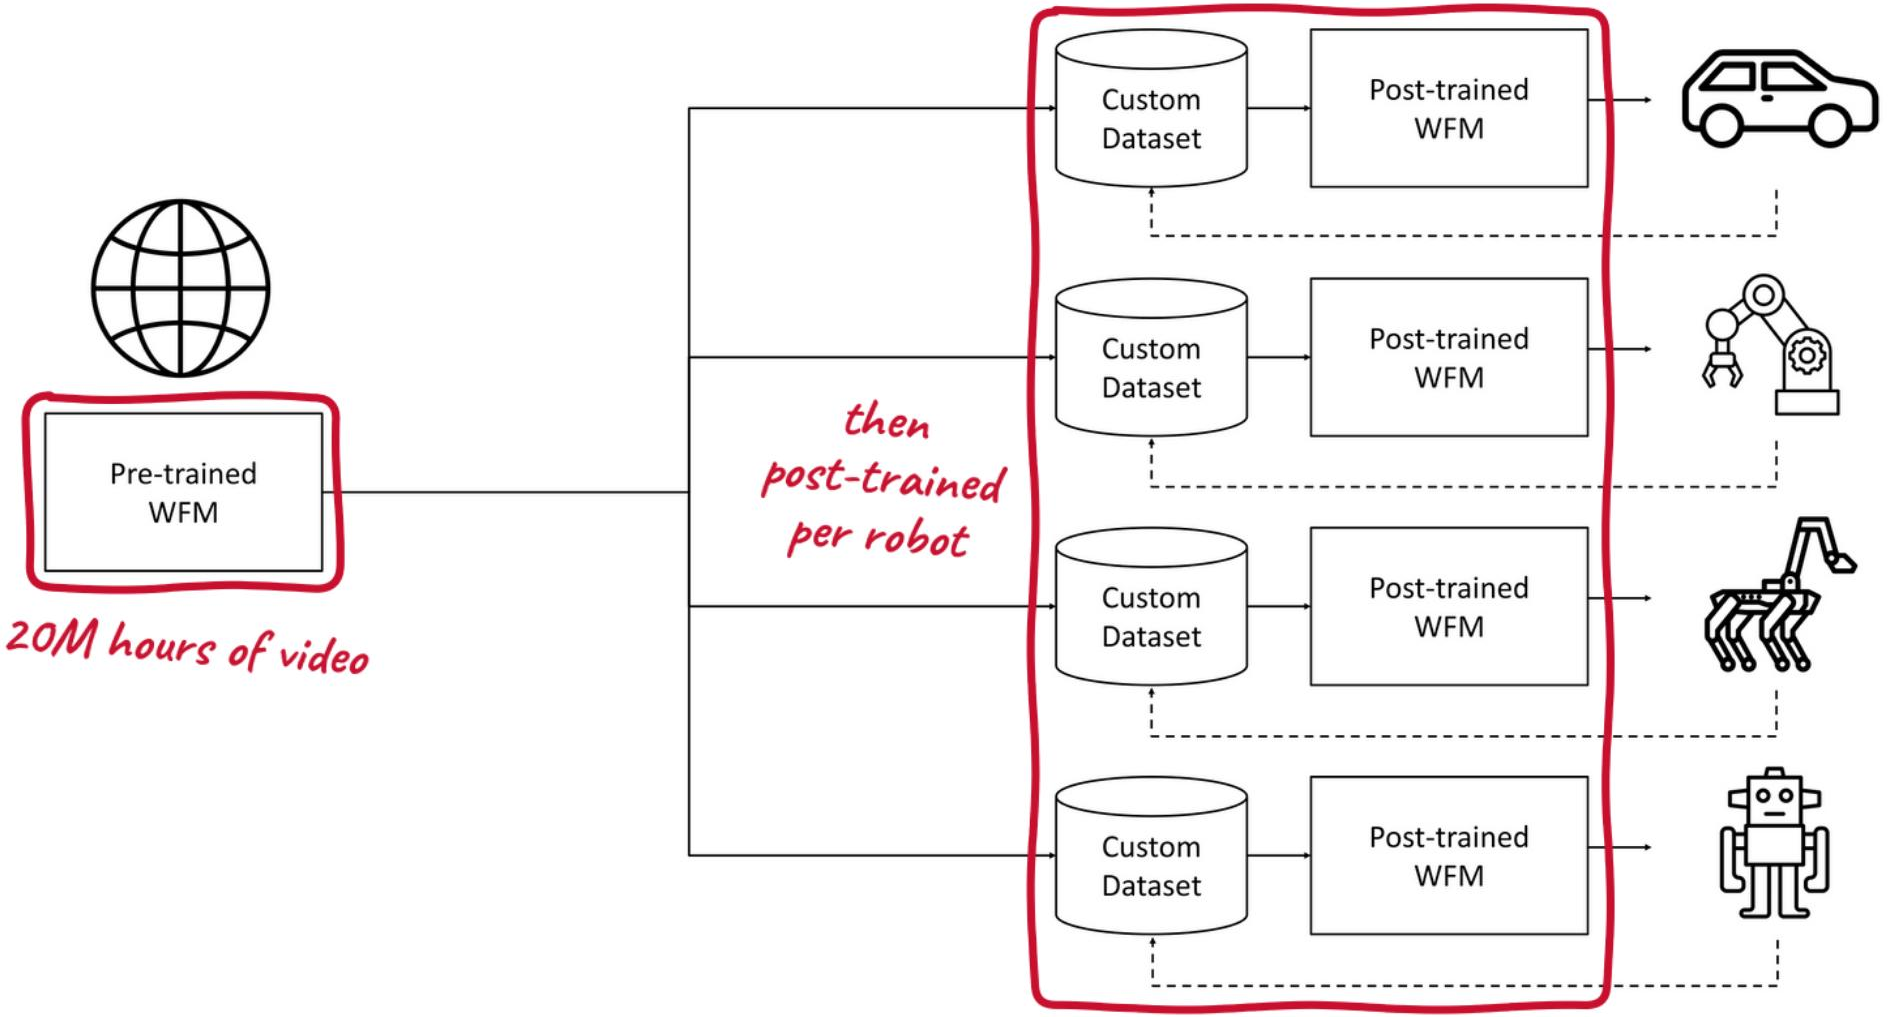
GAIA-1

Hu et al. 2023 · arXiv:2309.17080



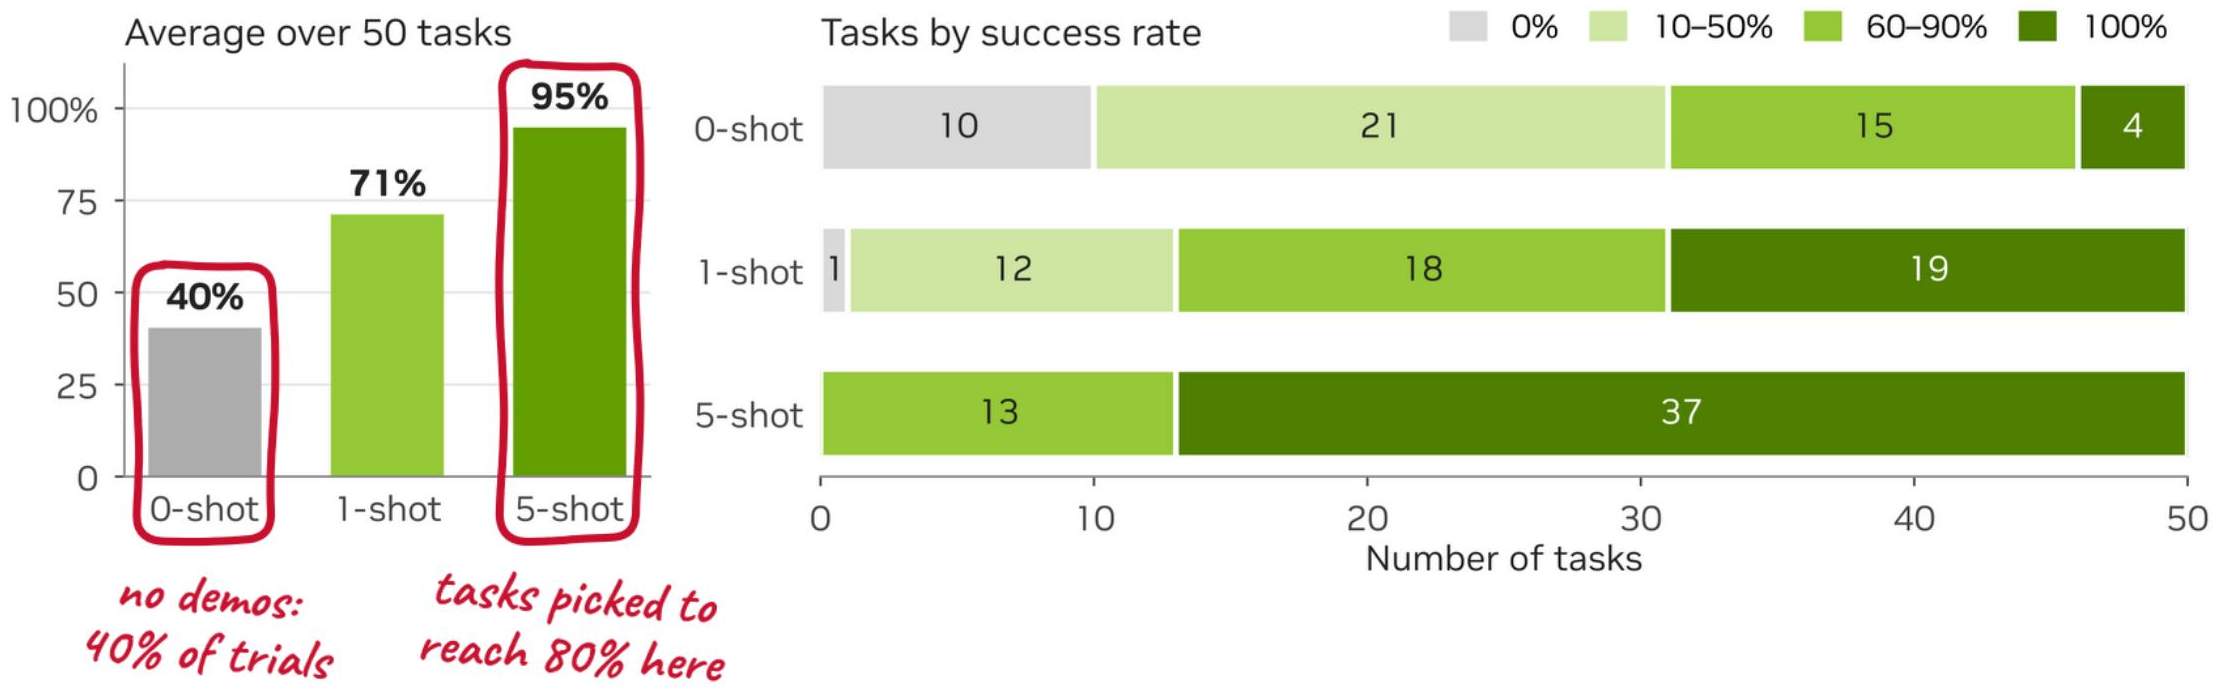
Cosmos

NVIDIA 2025 · arXiv:2501.03575



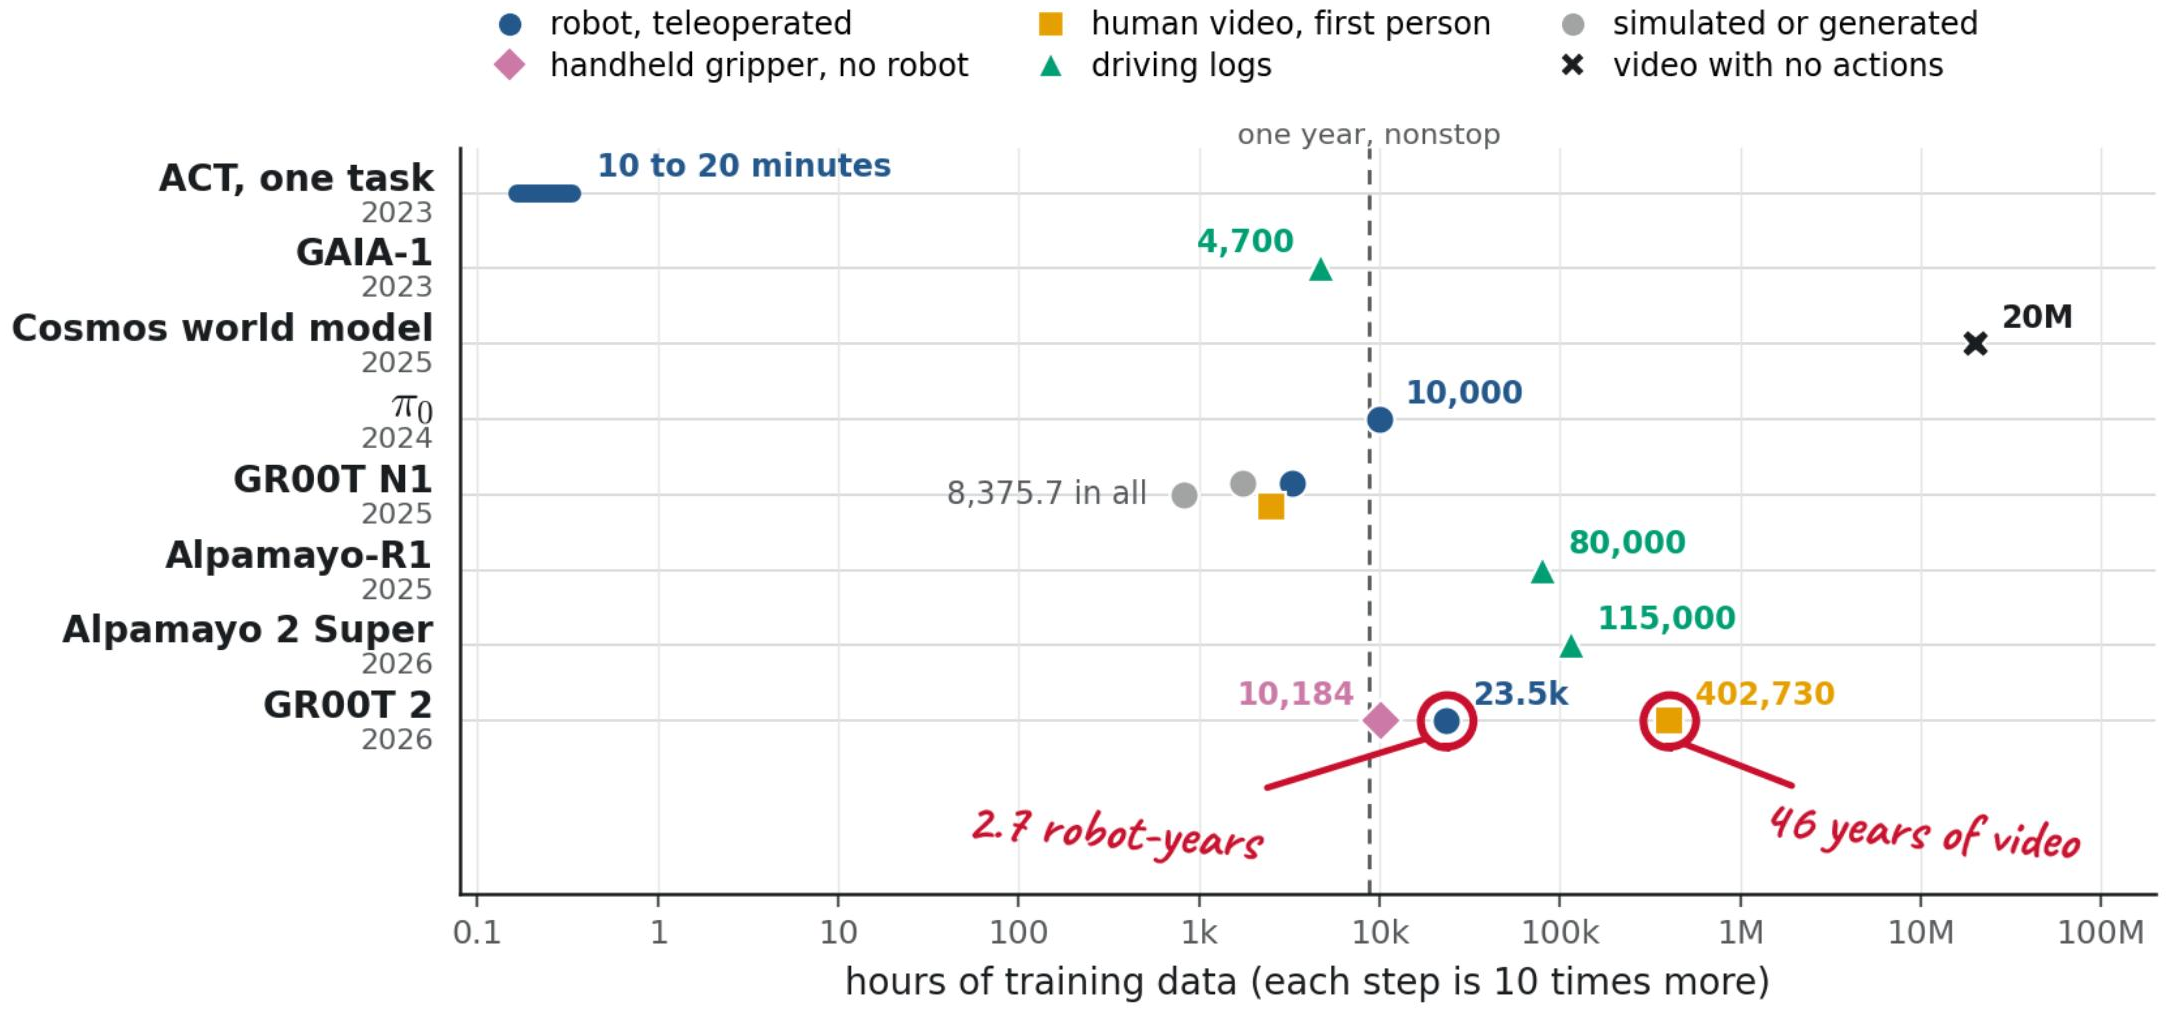
GR00T 2

NVIDIA 2026 · GR00T 2 whitepaper, November 9



Training Hours

hours as each paper or model card prints them



Reading

Required

Black et al., π_0: A Vision-Language-Action Flow Model for General Robot Control · arXiv:2410.24164

Optional

Brohan et al., RT-2 · arXiv:2307.15818 Kim et al., OpenVLA · arXiv:2406.09246

Chi et al., Diffusion Policy · arXiv:2303.04137 Zhao et al., ACT · arXiv:2304.13705

Hu et al., GAIA-1 · arXiv:2309.17080 NVIDIA, Cosmos · arXiv:2501.03575

Today's other systems

$\pi_{0.7}$ · arXiv:2604.15483 Alpamayo-R1 · arXiv:2511.00088 Cosmos-Reason1 · arXiv:2503.15558 GR00T N1 · arXiv:2503.14734

December 1 · last lecture **December 3** · poster session