From VLA to Agentic Manipulation

Yunlong Song

Guest lecture at ETH Zürich

November 23, 2026

Conference on Robot Learning (CoRL), 2026



"what do you think about Astra?"

Astra on a real-robot benchmark

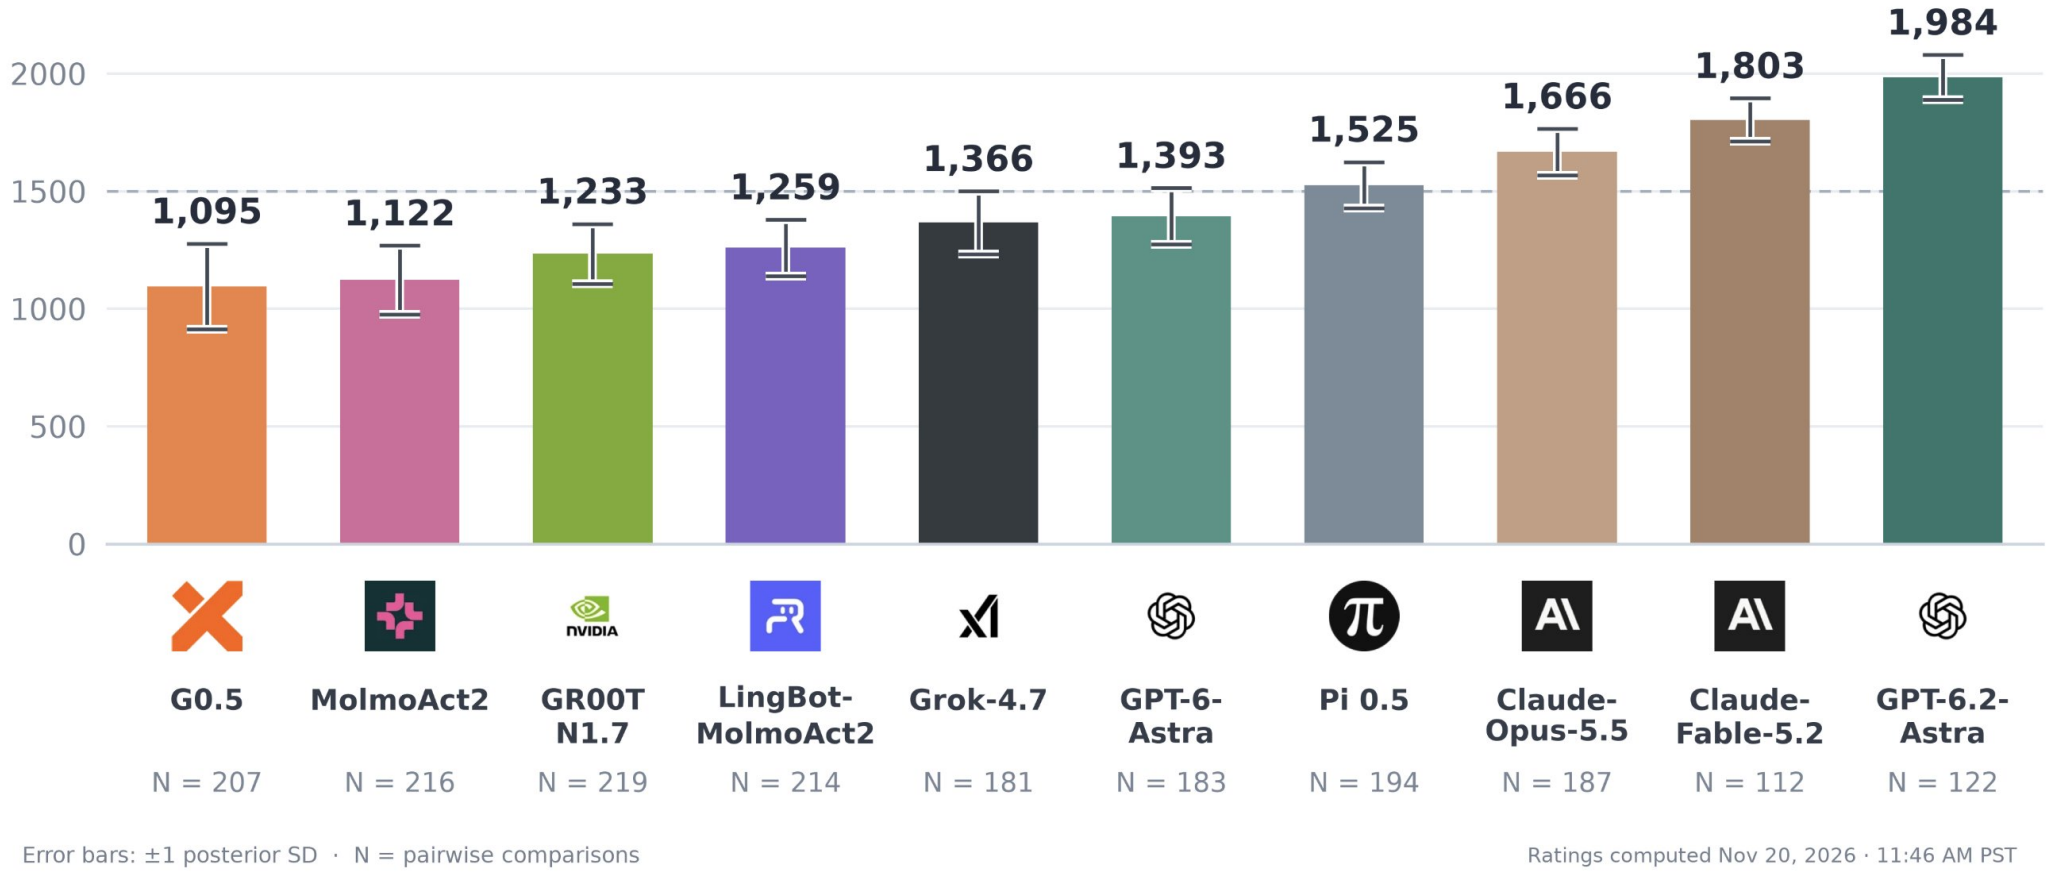

Source: RoboColosseum · x.com/RoboColosseum/status/2123974530859405631

Manipulation in the foundation-model era

Apr 2022
SayCan
Language models select robot skills
Ahn et al., 2022

Dec 2022
RT-1
Scale multi-task robot learning
Brohan et al., 2022

Jul 2023
RT-2
Transfer web knowledge into robot actions
Brohan et al., 2023

Oct 2023
Open X-Emb.
Pool data across robot embodiments
OXE Collab., 2023

Jun 2024
OpenVLA
Open models for generalist manipulation
Kim et al., 2024

Oct 2024
π0
VLM backbone with a flow action expert
Black et al., 2024

Apr 2025
π0.5
Generalize to unseen environments
Physical Intelligence

Nov 2025
π*0.6 / RECAP
Improve through experience and corrections
Physical Intelligence

Feb 2026
DreamZero
Joint video and action prediction
Ye et al., 2026

Nov 2026
AMP1
Generalist agents use tools and feedback
Chen et al., 2026

Note: RT-2 (Brohan et al., 2023) introduced the term "**VLA**".

Vision-language-action models (VLA)

Vision-language-action models

Learn a policy that maps images, instructions, and robot state to robot commands.

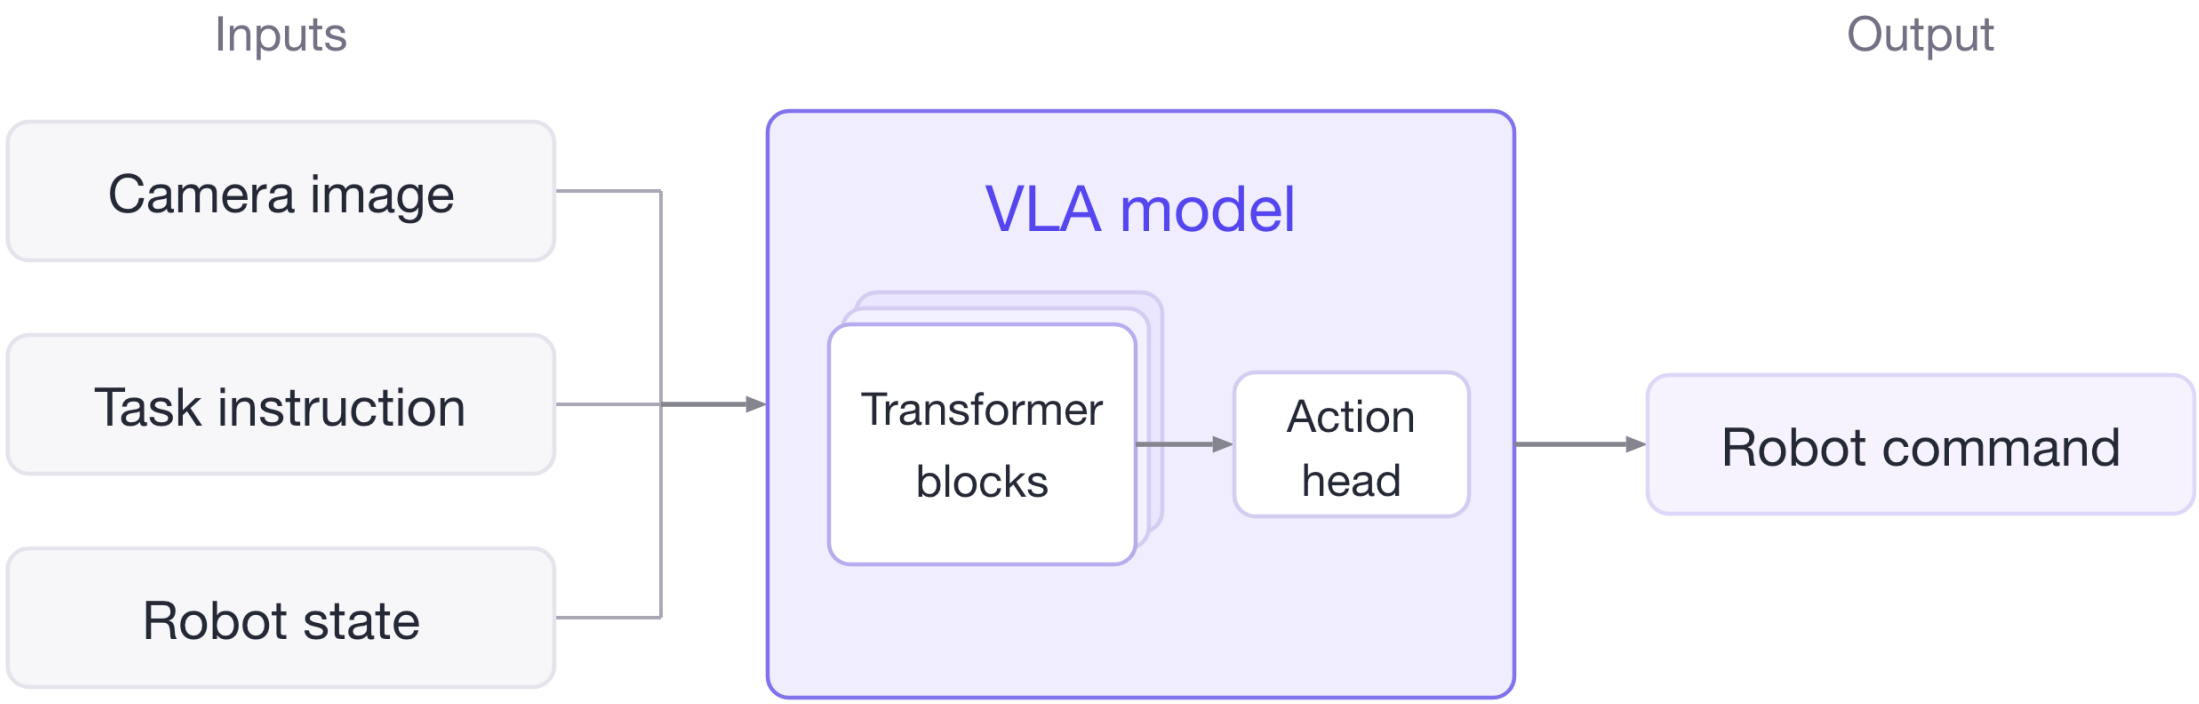

RT-2: knowledge transfer

Transfer knowledge from images and language into robot actions.



Start with a VLM

Train on robot actions

Predict commands

Demo instruction: “pick up the extinct animal.”

π0: action generation

Generate continuous robot actions with a VLM and an action expert.

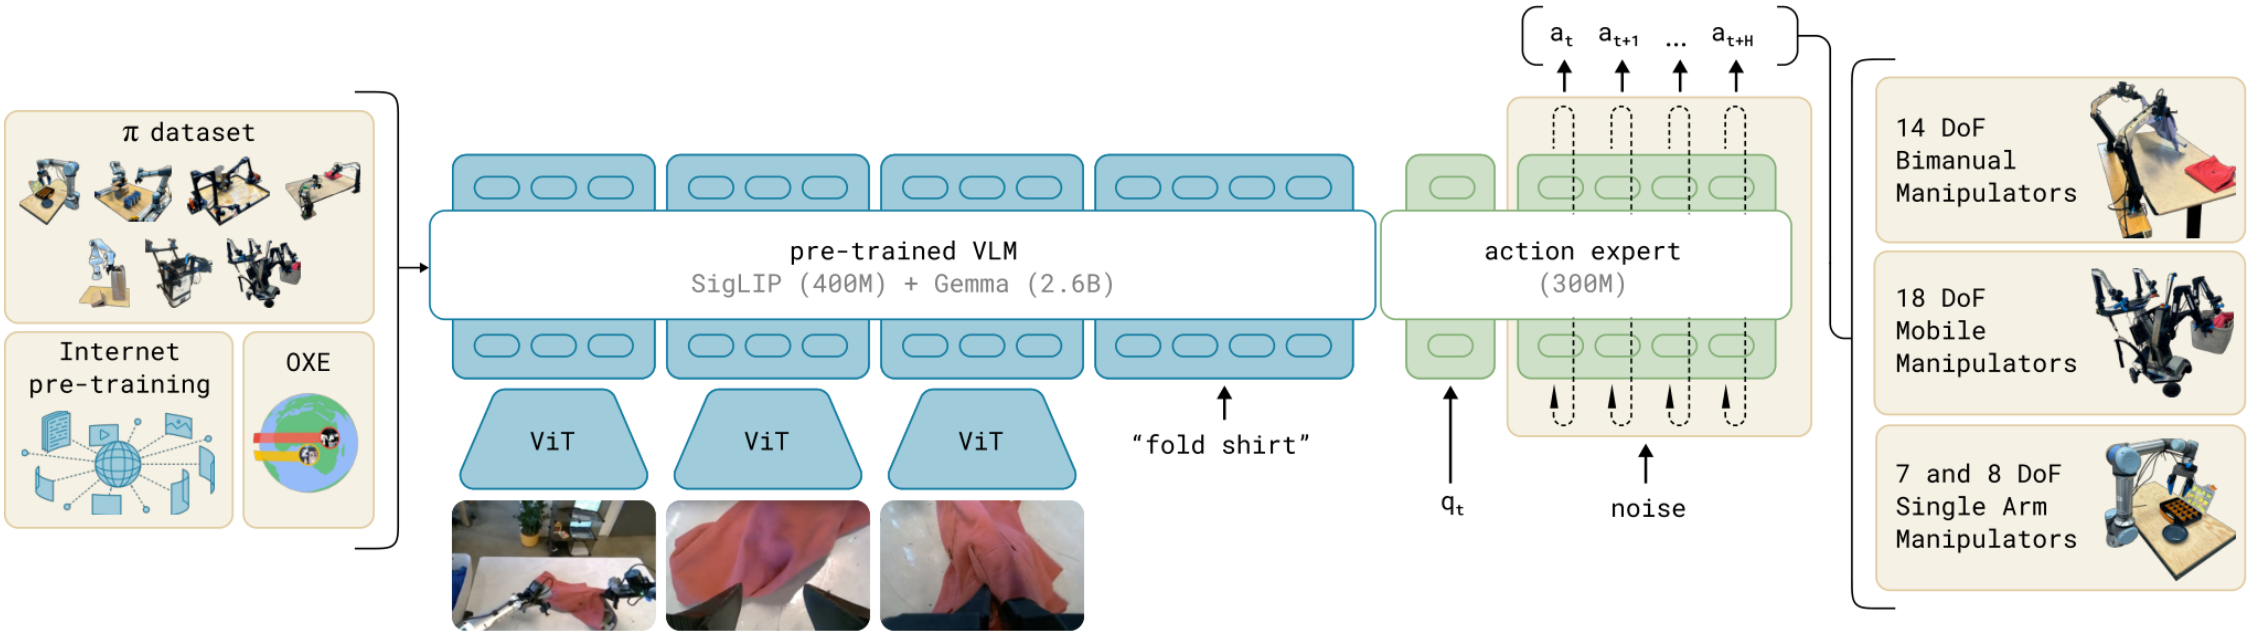

Model

VLM backbone
+ action expert.

Actions

Flow matching generates
continuous action chunks.

Training

Learn across robots,
then specialize for tasks.

π0.5: generalization

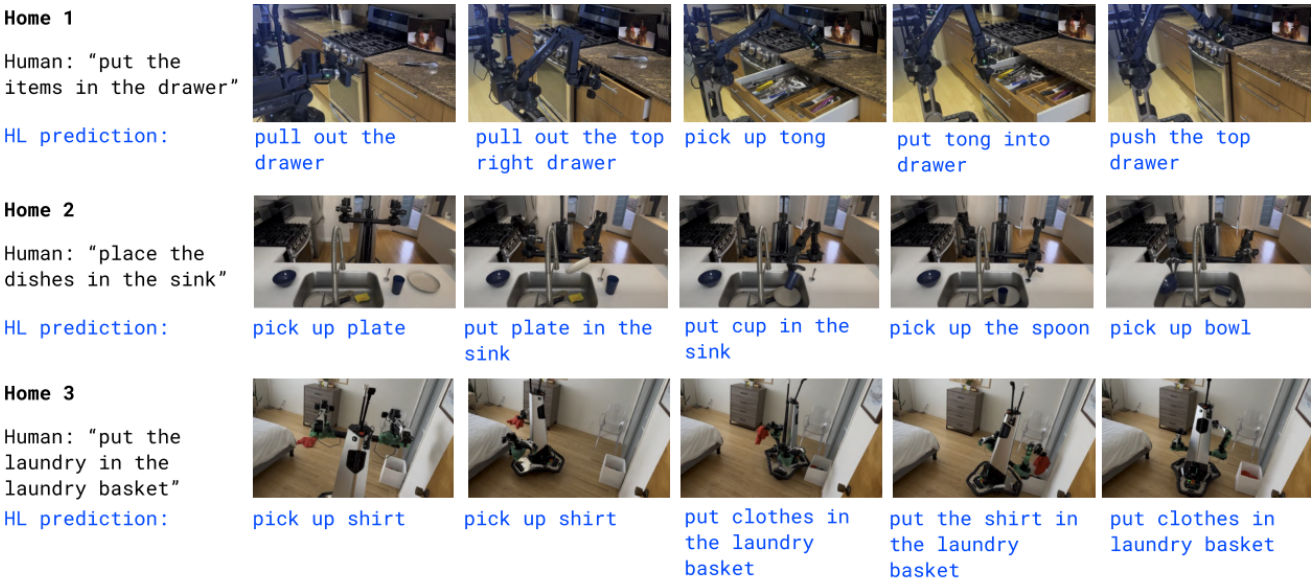

Diverse data

Combine robot demonstrations, web data, and subtask labels.

Subtasks before actions

Predict what to do next, then generate the motion.

π*0.6 / RECAP: learning from experience

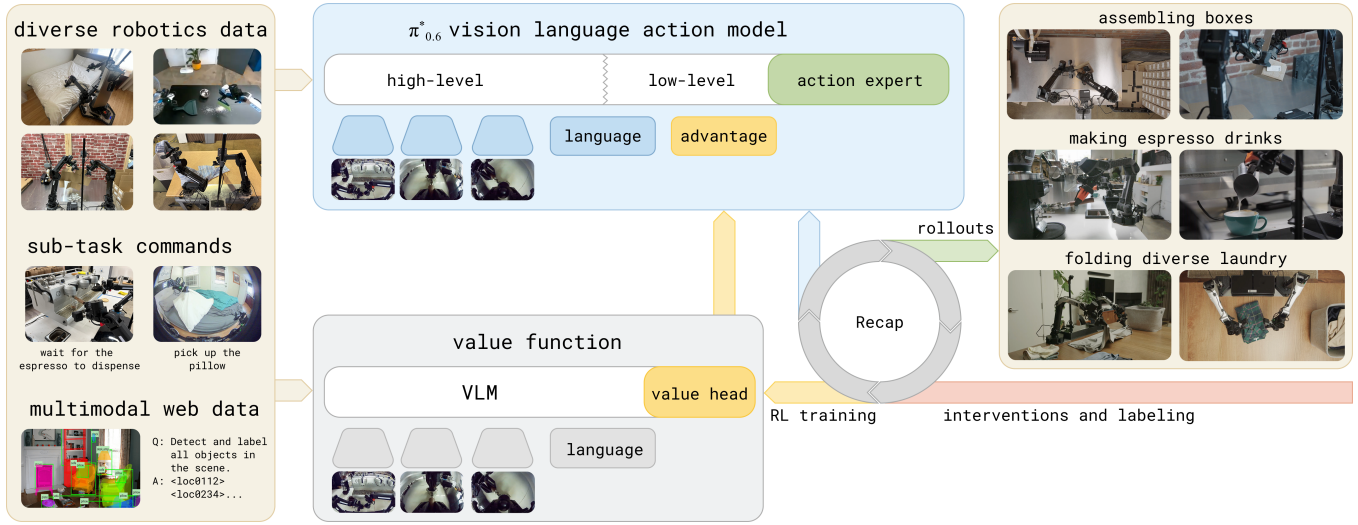

Collect experience

Estimate progress

Update the policy

RECAP: the training method.

π*0.6: the resulting policy.

RoboTTT: reasoning and memory

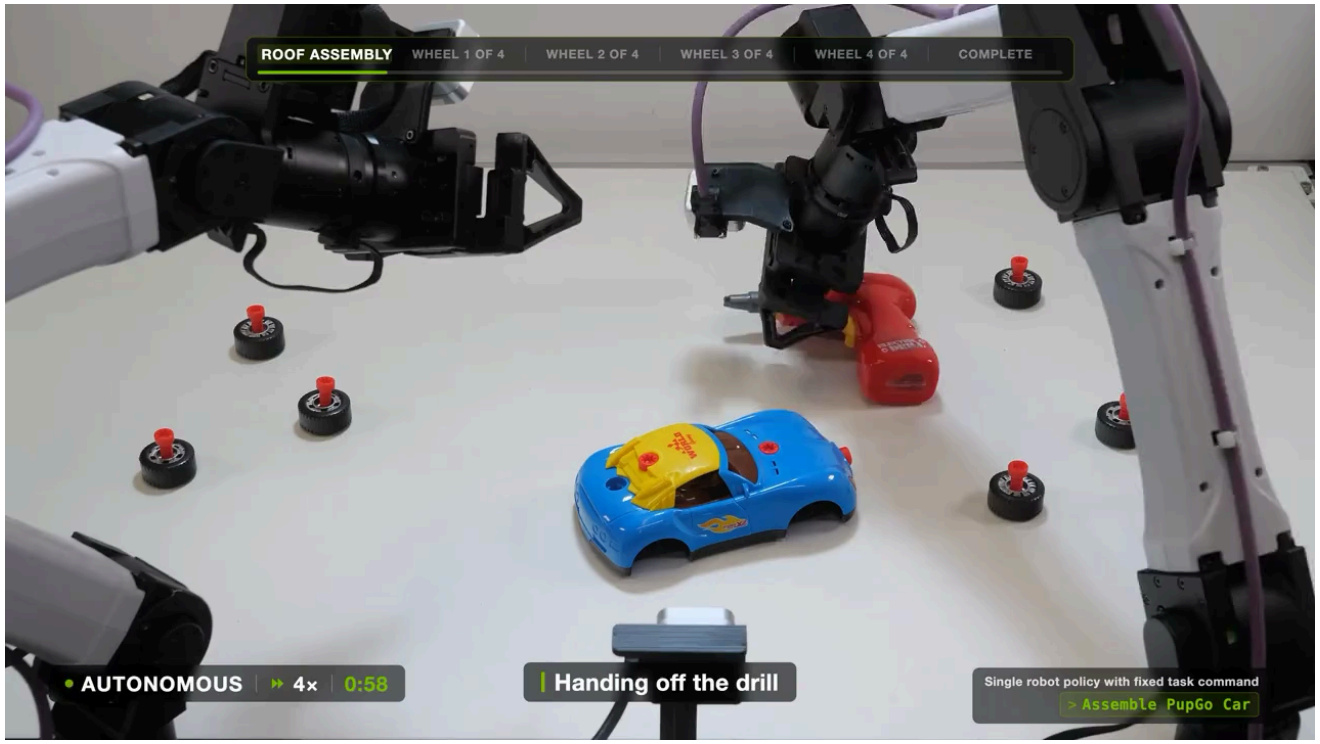

Demo: toy-car assembly with memory.

Remember

Store past observations and actions in a compact memory.

Update during execution

Update the memory as the robot acts.

Longer tasks

Demonstrated five-minute, ten-stage robot assembly.

Jiang et al., 2026 · arXiv:2607.15275 · 35 s excerpt; source video at 4× speed

Common VLA challenges

- Generalizing to new tasks remains difficult.
- Learning actions requires substantial robot demonstration data.
- Videos without action labels are harder to use for action training.

DreamZero (Ye et al., 2026), arXiv:2602.15922; NVIDIA WAM overview (Reuss, 2026).

World-action models (WAM)

What is a world-action model?

A robot policy that uses learned world dynamics to generate actions.

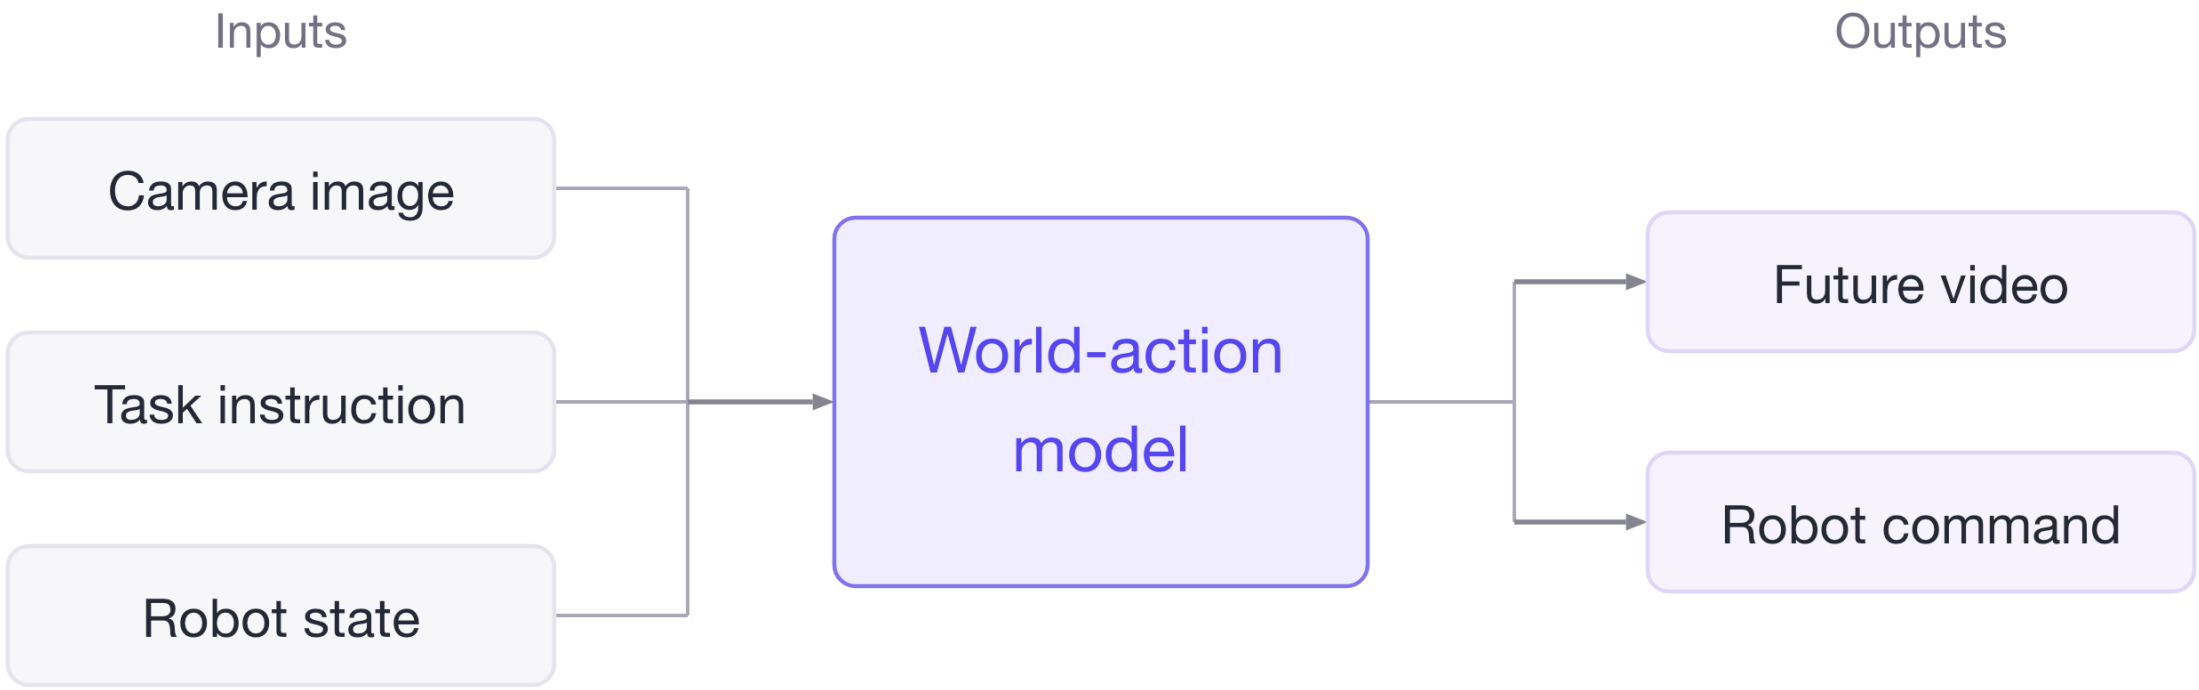

Example: DreamZero jointly predicts future video and robot commands.

How WAMs can help?

- Learn motion and interaction patterns from video.
- Connect predicted future states to robot actions.
- Use (internet scale) video-only data alongside robot demonstrations.

DreamZero (Ye et al., 2026), arXiv:2602.15922; NVIDIA WAM overview (Reuss, 2026).

Three ways to use world-model knowledge

UniPi

Video first

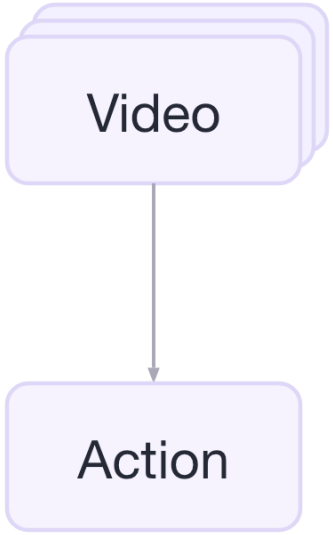

DreamZero

Joint prediction

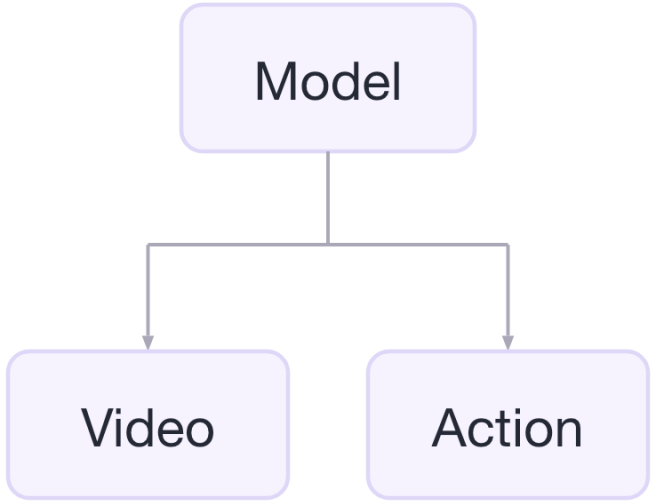

Fast-WAM

Features only

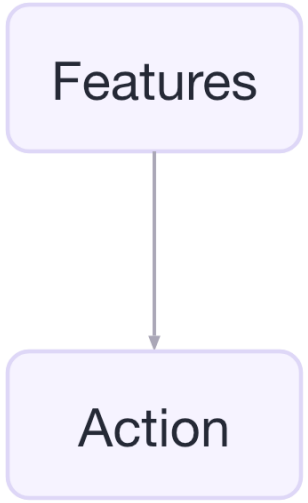

DreamZero: prediction and execution

Seen-task example: pick up the banana and place it on the plate.

Predicted future

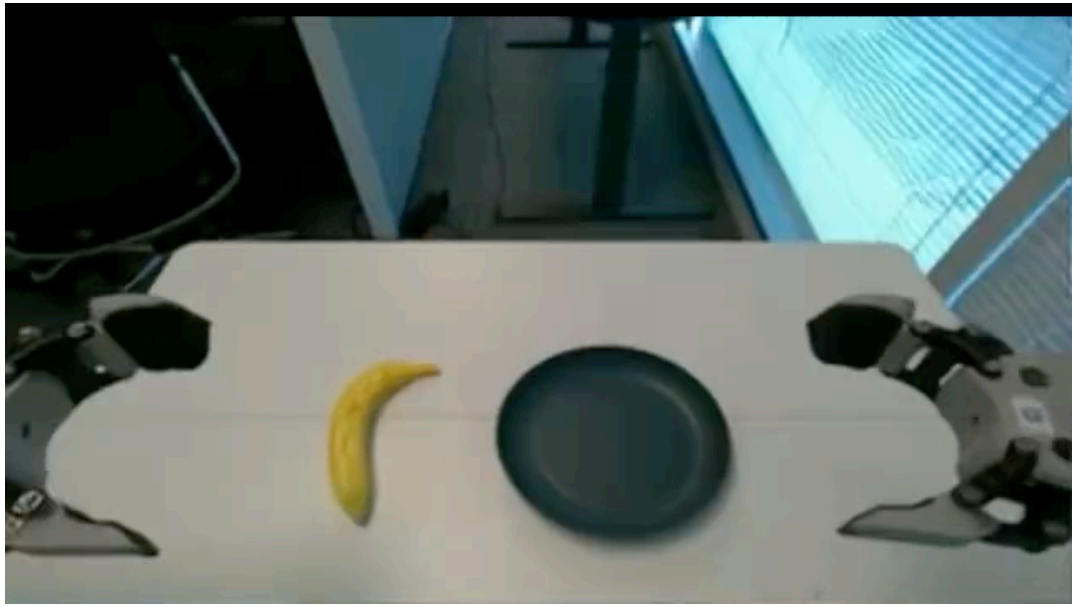

Generated video, robot-camera view

Actual robot execution

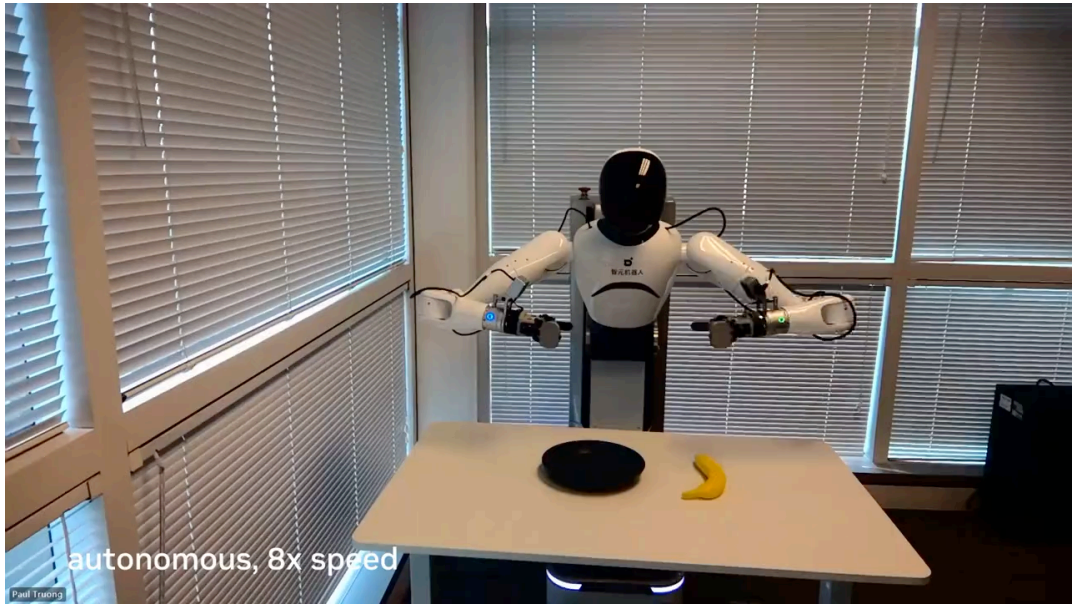

External camera, source at 8× speed

DreamZero (Ye et al., 2026). Same trial, different views and playback speeds.

Agentic manipulation

Learned policies vs. agentic policies

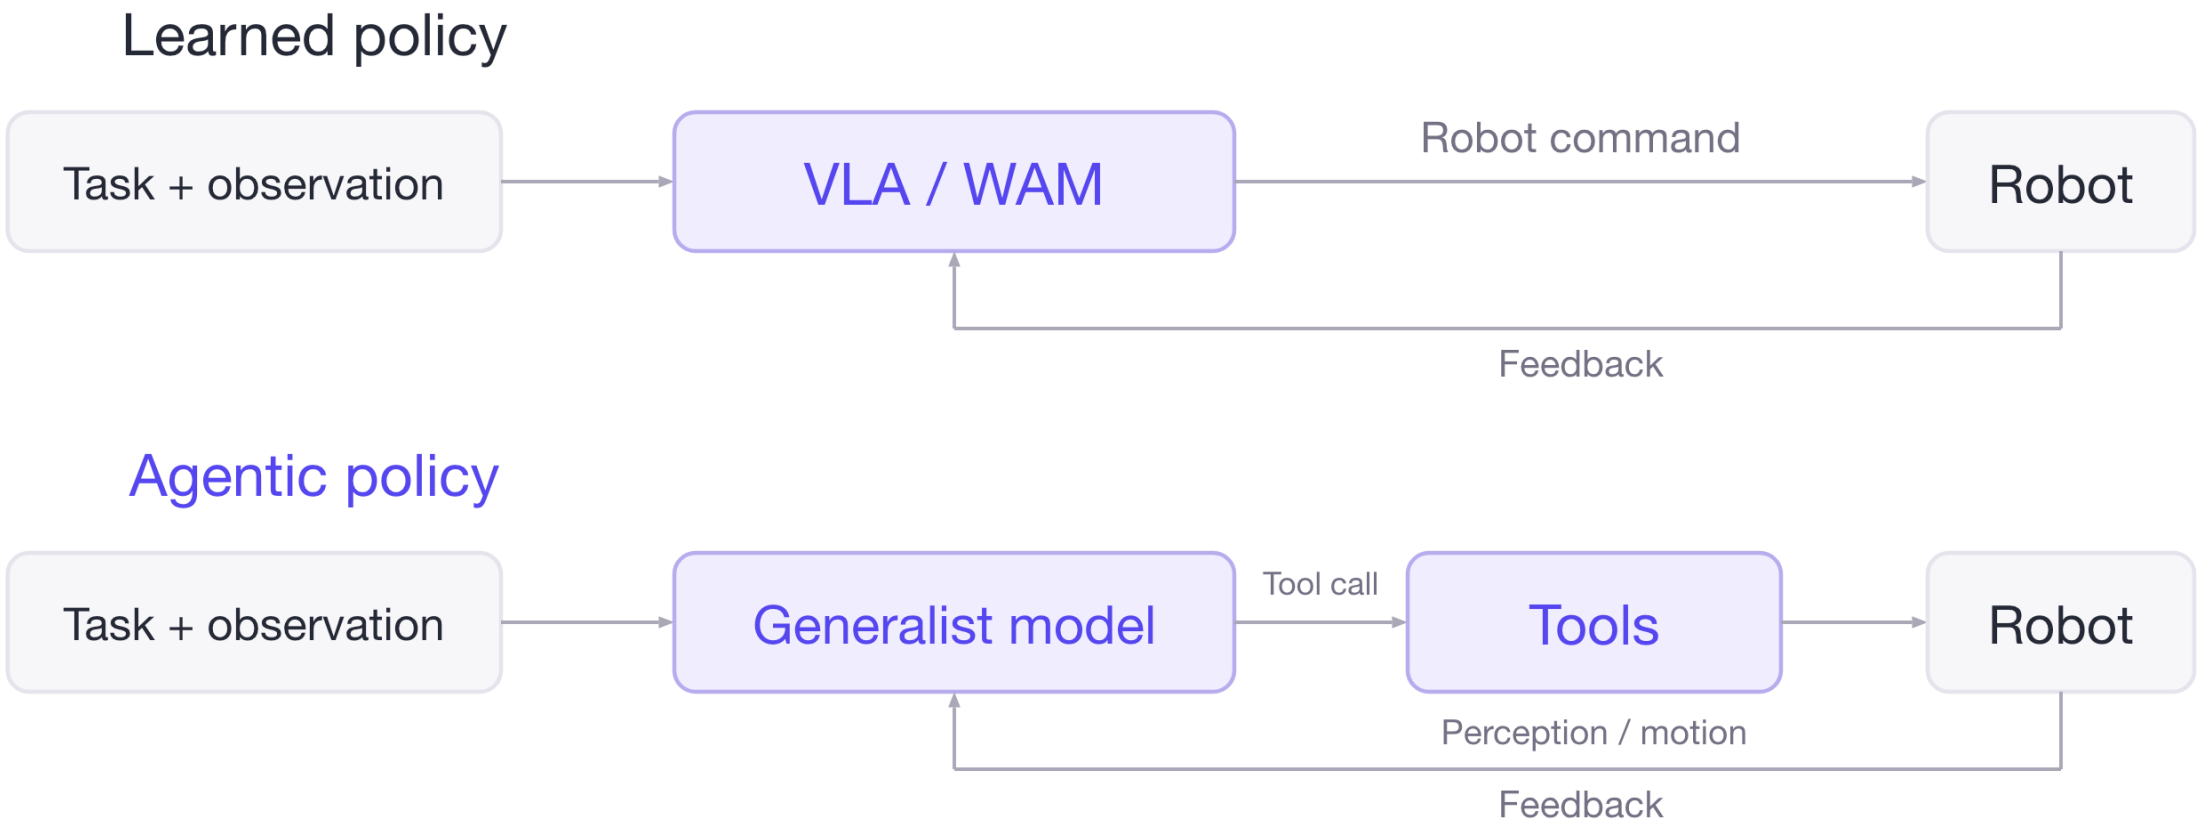

Astra and robot manipulation

Can a general-purpose model become part of a robot's policy?

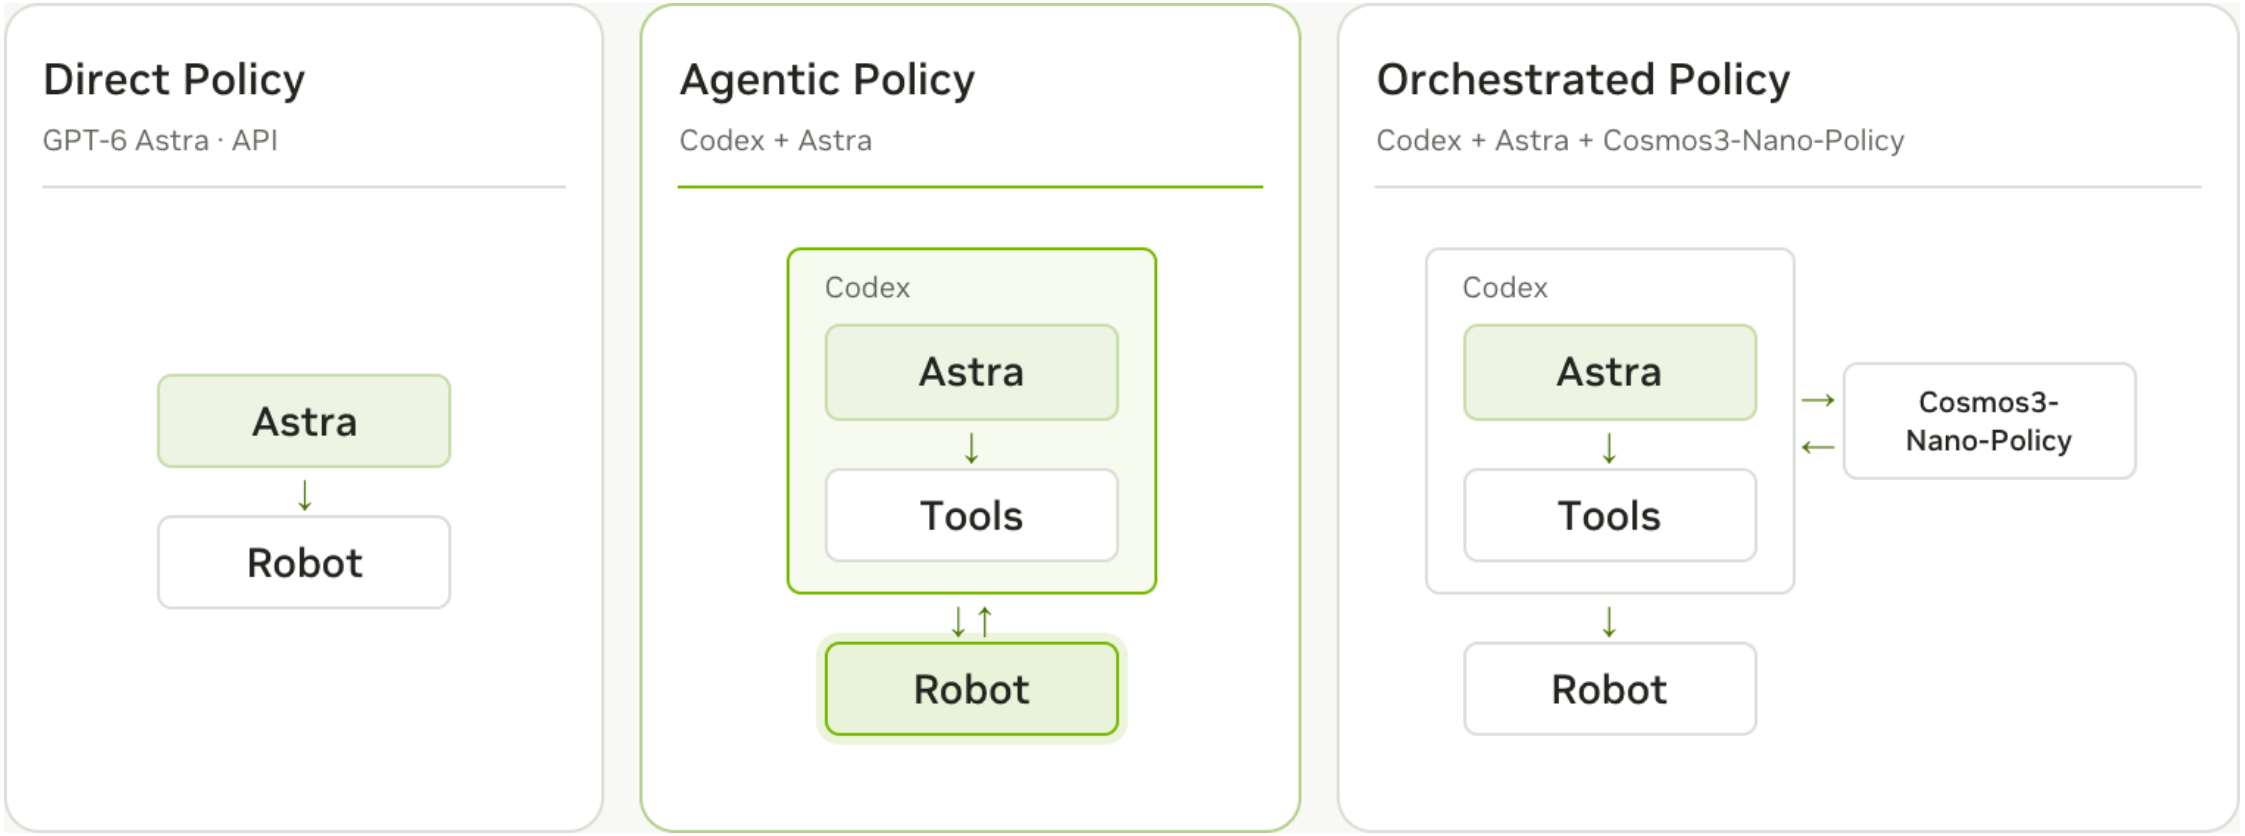

Yang et al., Astranomy (2026), Fig. 4 · research.nvidia.com/labs/srl/projects/astranomy/

Astranomy: general-purpose robot control

GPT-6 Astra outputs robot commands without robot-policy fine-tuning.

RoboLab-120 simulation · success rate

Direct policy

66.3%

Astra API

Agentic policy

74.8%

Astra + history and feedback

Best reported in this study

Orchestrated policy

66.5%

Astra + robot policy

Yang et al., Astranomy (2026) · research.nvidia.com/labs/srl/projects/astranomy/

Astranomy: direct robot control

Task: put hammers in the right bin without touching other objects.

Direct policy: Astra generates robot commands through an API.

Left shoulder camera | Right shoulder camera | Wrist camera

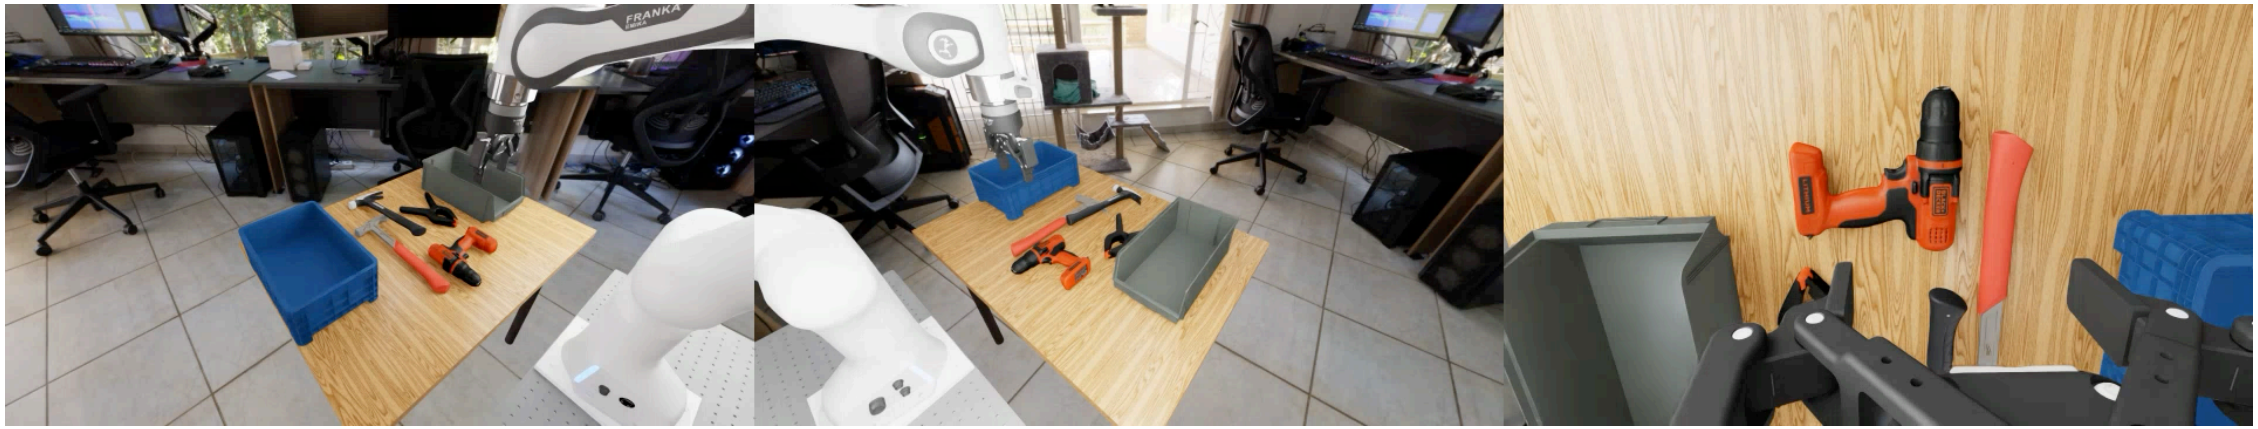

On Agentic Manipulation Policies (Chen et al., 2026)

Task: pick up the banana and place it in the bowl.



1. Observe and plan

TOOL CALLS

```
observe()
objects()
preview()
```

FEEDBACK

Locate banana and bowl.
Check the approach path.

Summary of the official AMP1 trace

Emergent behaviors in AMP1

Selected real-robot demos · 20× playback; reasoning time omitted.

Search for hidden objects

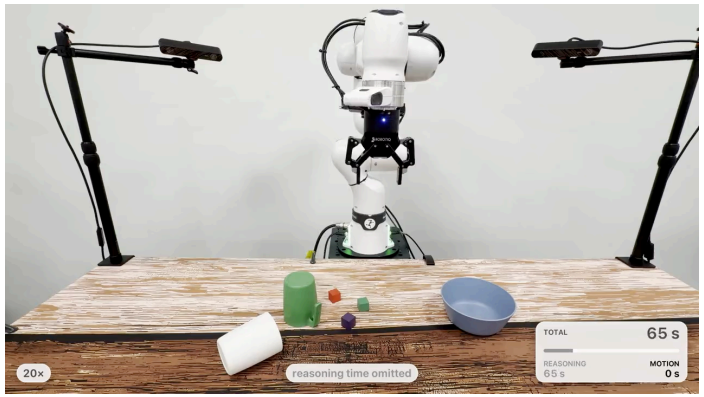

Moves a covering cup,
then retrieves the yellow block.

Make an object graspable



Repositions the book,
then grasps it.

Use a physical tool



Uses a spoon to bring the
bottle within reach.

Chen et al., On Agentic Manipulation Policies (2026), Section 2B. Videos: official AMP1 project.

Key takeaways

- VLA: maps visual observations and instructions to robot commands.
- WAM: uses learned world dynamics to support action generation.
- Agentic manipulation: uses a generalist model to call tools, inspect feedback, and revise actions.
- These approaches can complement one another in the same system.
- The key challenge: combine flexible reasoning with fast, reliable control.

Where does robotics go next?

A full-stack approach

Data	Capture physical skills and real-world experience
Model	Learn capabilities that transfer across tasks
Evaluation	Test reliability and expose failures
Hardware	Turn intelligence into useful physical work

Genesis AI

Human-skill data



Capture human motion and touch to teach robots physical skills.

GENE



Learn dexterous skills from demonstrations with a shared model.

Genesis World

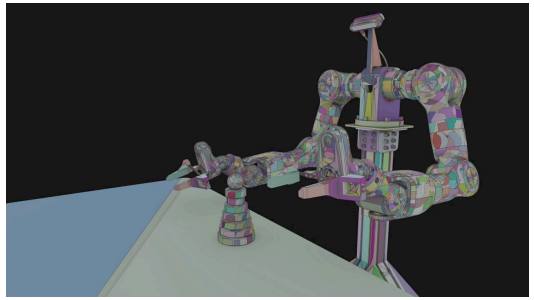

Simulate scenarios to test and improve robot behavior.

Eno

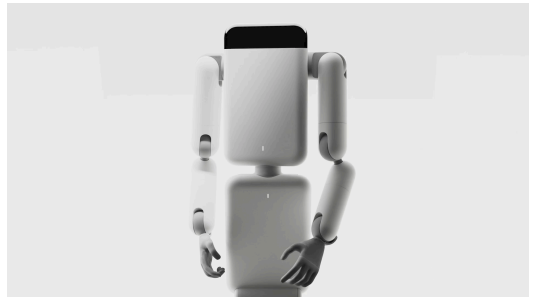

Bring intelligence into the real world through physical work.



Now is the most exciting time to do robot manipulation.

We are hiring

Internship/Full-time

Simulation, Rendering, Model Training,
Data, Robot Systems, Robot Learning

Bay Area, US · Paris, France · London, UK · Hybrid

careers@genesis.ai