Why General-Purpose Robots Cannot Be Pretrained

From Pretraining to Experience-Scaling Embodied AI

Dr. Jun Jin

jjin5@ualberta.ca

Assistant Professor
Department of Electrical and Computer Engineering
University of Alberta
Amii Fellow

- Key Words of my research:

Robotics

Mobile Robots

Continual learning

Obstacle Avoidance

Skill Discovery

Autonomous Navigation

Reinforcement learning

Language informed RL

Robotic Manipulation

Embodied AI

Vision-guided Robot Control

Learning from demonstration

Visual Servoing

Surgical robots

Diffusion Models

Dr. Jun Jin

Assistant Professor
Department of Electrical and Computer Engineering
University of Alberta
Amii Fellow

Website
Google Scholar

The Rise of Embodied AI

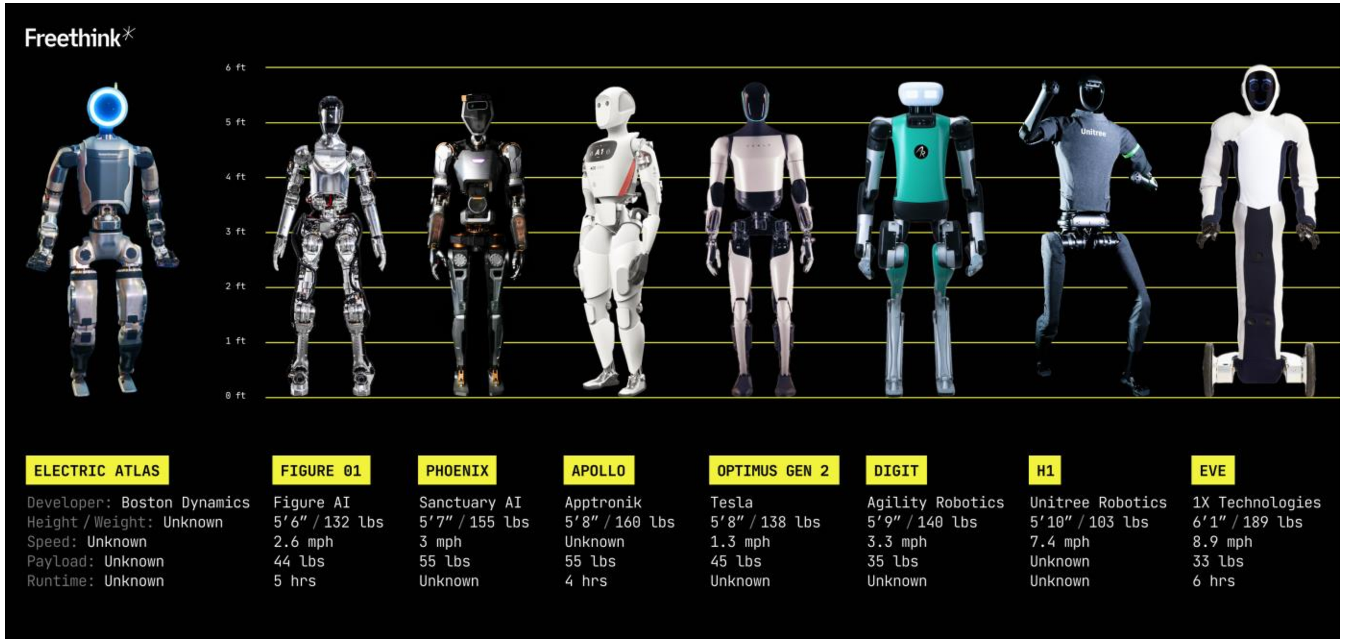

https://www.freethink.com/robots-ai/humanoids

The Rise of Embodied AI

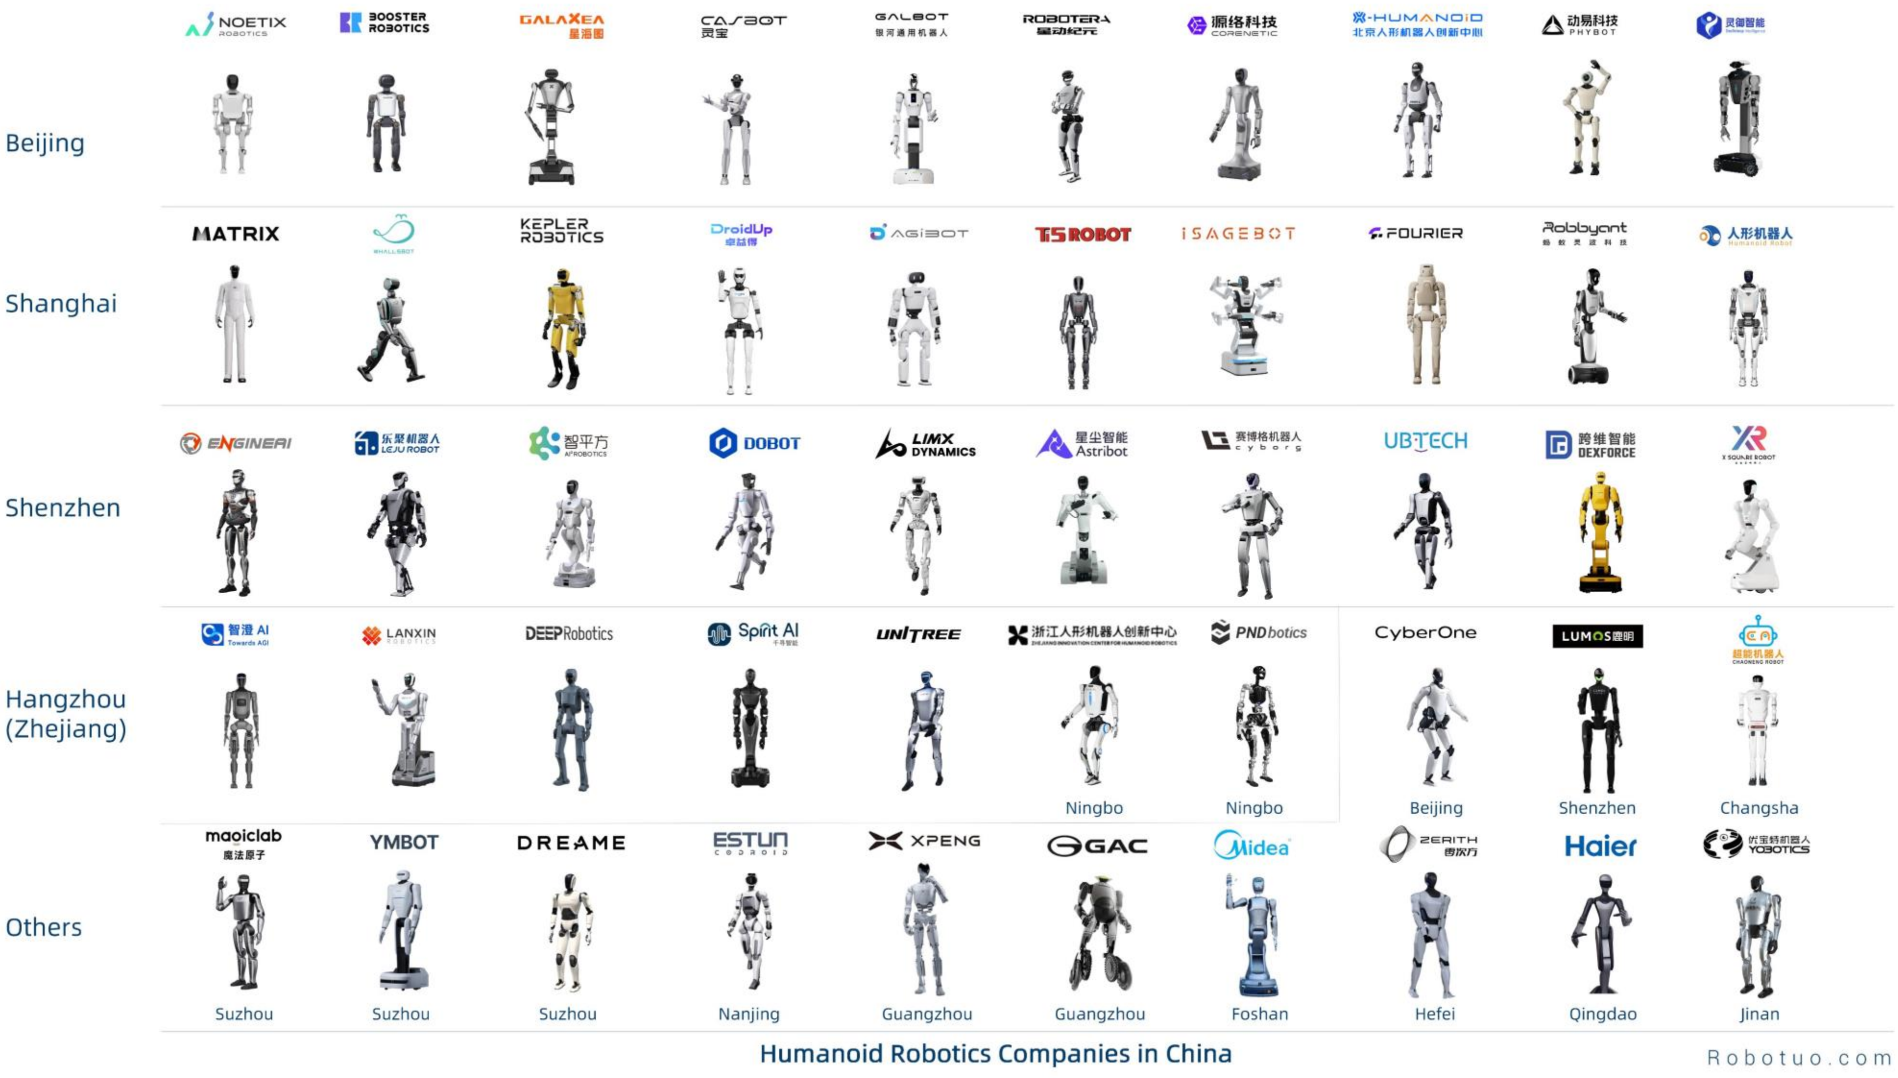

https://www.freethink.com/robots-https://hardwarefyi.substack.com/p/humanoid-robotics-mappedai/humanoids

The Rise of Embodied AI

KEPLER ROBOTICS
UBTECH
UNITREE
mentee robotics
1X
Boston Dynamics
FIGURE
FOURIER
SANCTUARY AI

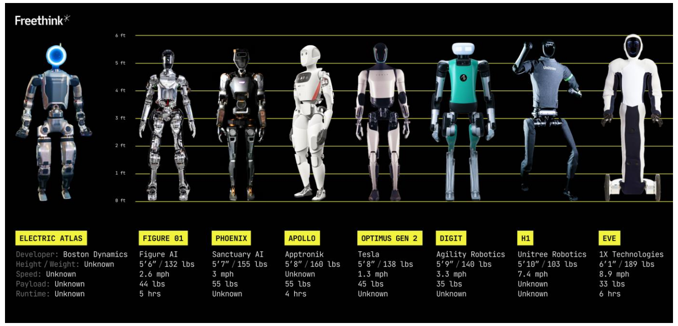

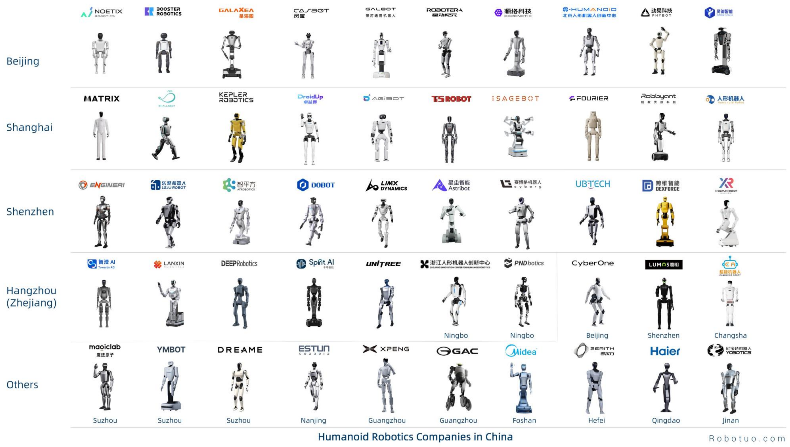

Humanoid Robotics Companies in China

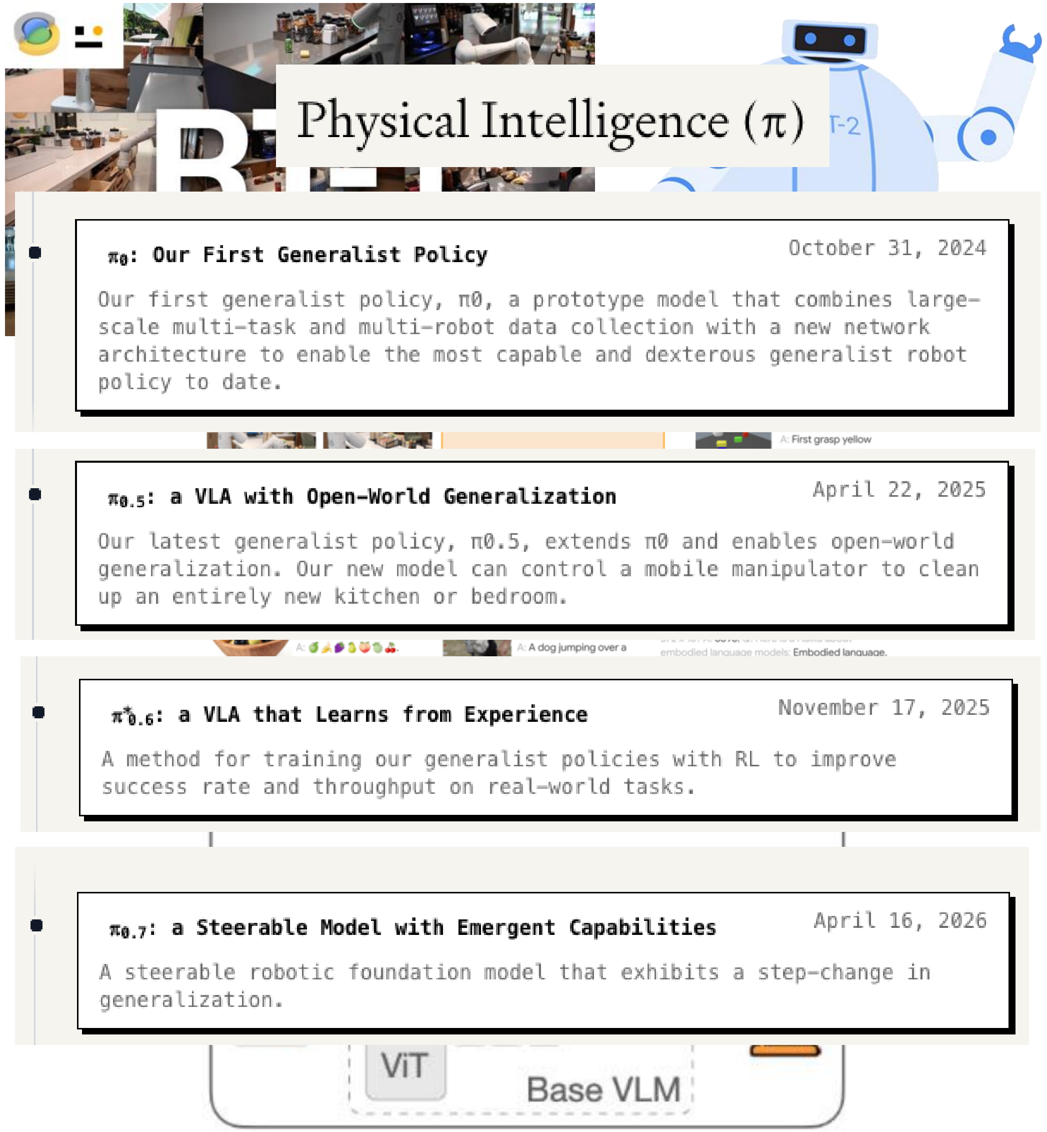

https://robotics-transformer1.github.io/ https://palm-e.github.io/ https://www.pi.website

https://robotics-transformer2.github.io

What is Embodied AI (EAI) about?

The Pursuit of General-Purpose Robots

I have different views, will revisit this point later.

Two robots in a kitchen: one answering a human's question about cereal, the other washing dishes. Scene showcases AI skills in navigation, manipulation, and reasoning.

Illustration courtesy of Winson Han.

https://engineering.oregonstate.edu/all-stories/beyond-chatgpt-embodied-language-understanding

The prevailing Bet in Embodied AI

➢Robotics will follow the same pipeline shape as LLMs:

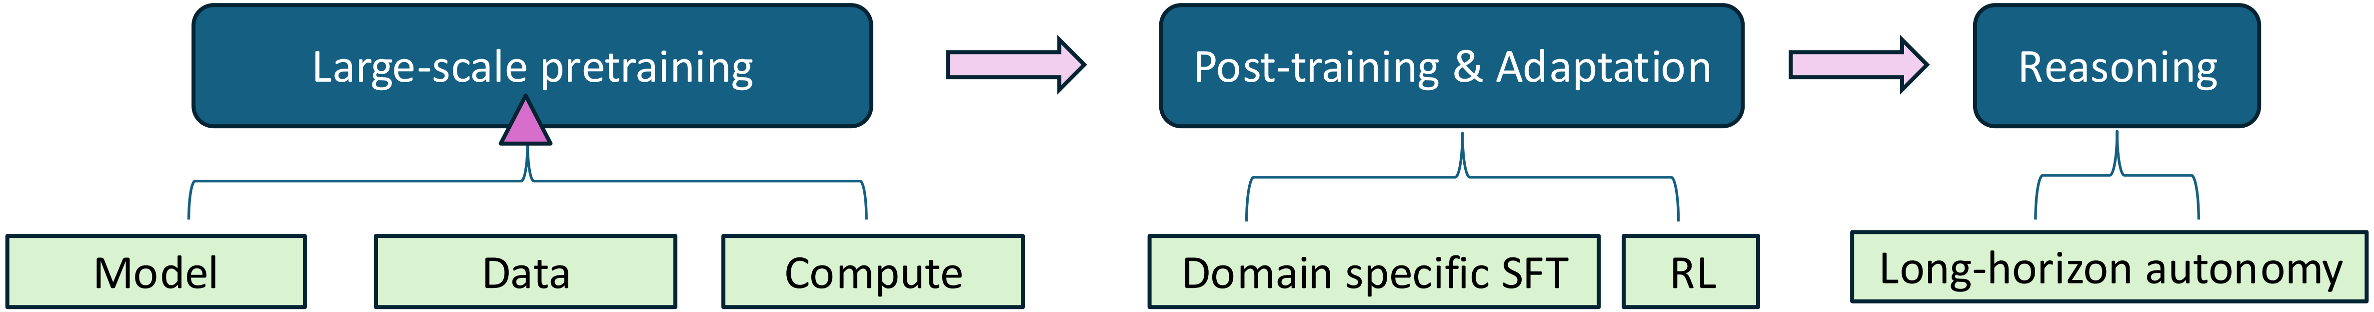

- **(1) The prevailing expectation is that larger models, larger robot-relevant datasets, and greater compute will yield broader robotic task capability.**

with the ultimate goal

General-Purpose Robots

The prevailing Bet in Embodied AI

➢**Robotics will follow the same pipeline shape as LLMs:**

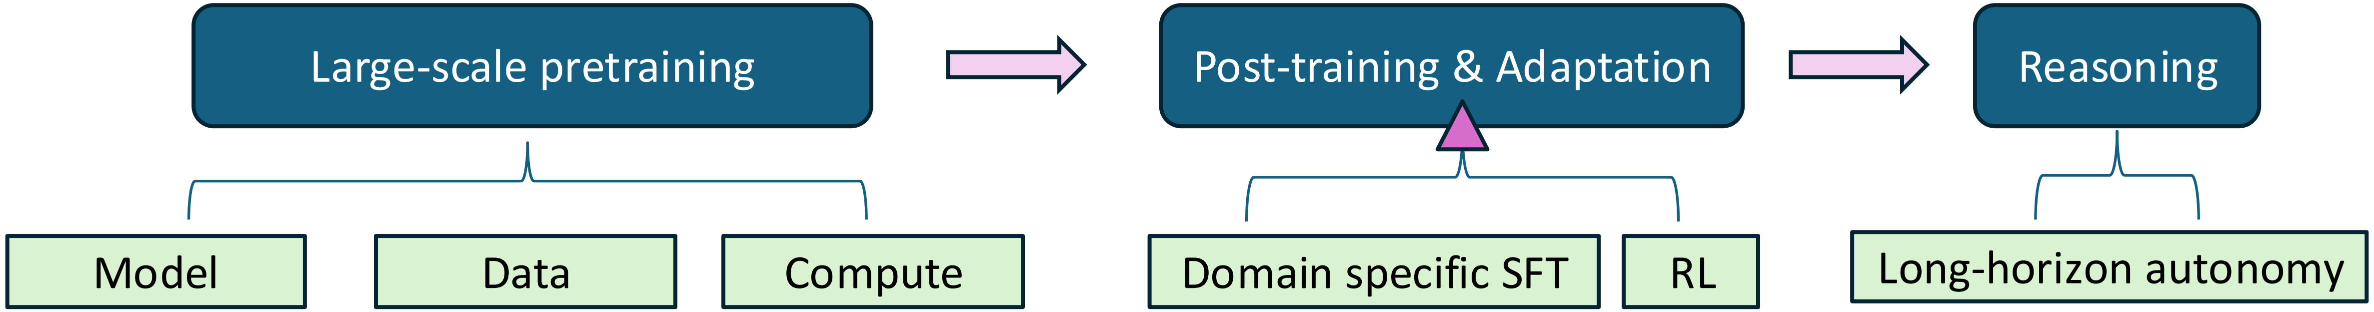

- (1) The prevailing expectation is that larger models, larger robot-relevant datasets, and greater compute will yield broader robotic task capability.
- **(2) With a sufficiently strong pretrained backbone, domain-specific post-training (SFT, RL), is expected to become easier and more data-efficient.**

The prevailing Bet in Embodied AI

➢**Robotics will follow the same pipeline shape as LLMs:**

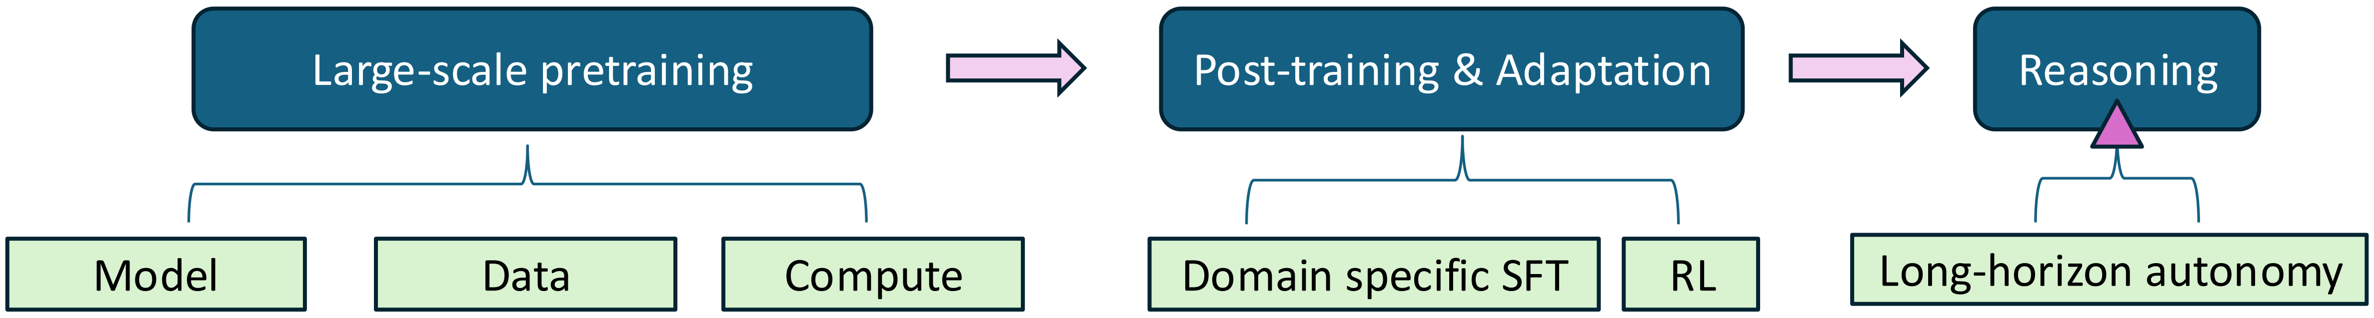

- (1) The prevailing expectation is that larger models, larger robot-relevant datasets, and greater compute will yield broader robotic task capability.
- (2) With a sufficiently strong pretrained backbone, domain-specific post-training (SFT, RL), is expected to become easier and more data-efficient.
- **(3) Keep scaling the backbone will expect to improve embodied reasoning and thus support long-horizon autonomy.**

The prevailing Bet in Embodied AI

➢**Robotics will follow the same pipeline shape as LLMs:**

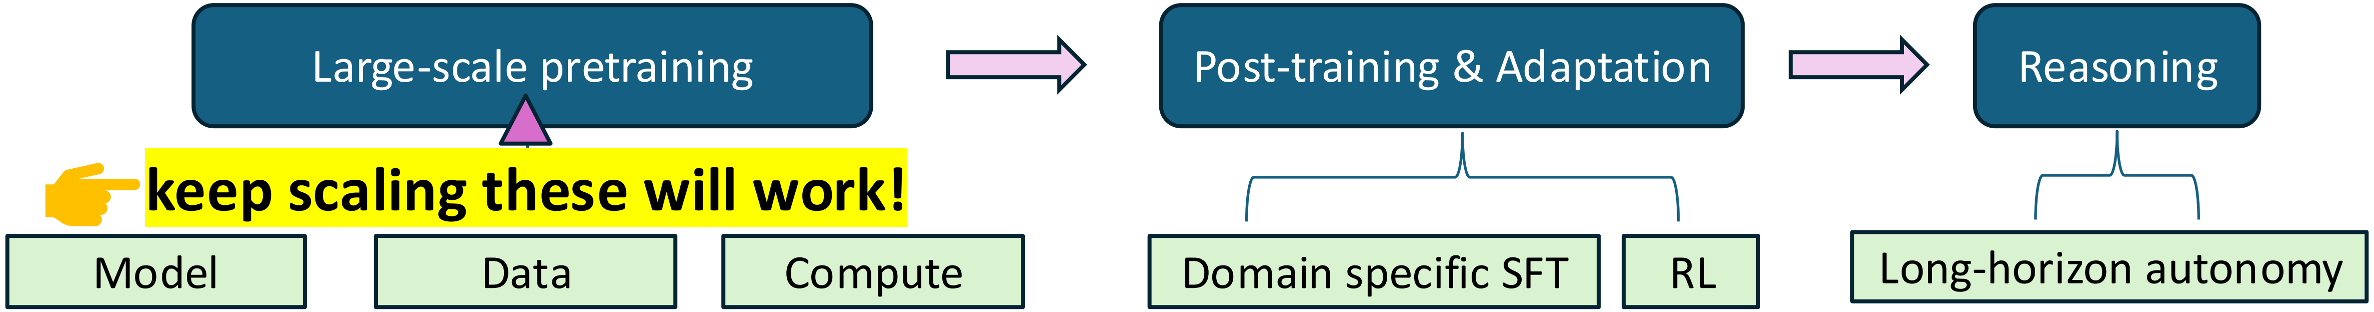

- (1) The prevailing expectation is that larger models, larger robot-relevant datasets, and greater compute will yield broader robotic task capability.
- (2) With a sufficiently strong pretrained backbone, domain-specific post-training (SFT, RL), is expected to become easier and more data-efficient.
- (3) Keep scaling the backbone will expect to improve embodied reasoning and thus support long-horizon autonomy.

The prevailing Bet in Embodied AI

➢Why the LLM Pipeline Matters for Robotics?

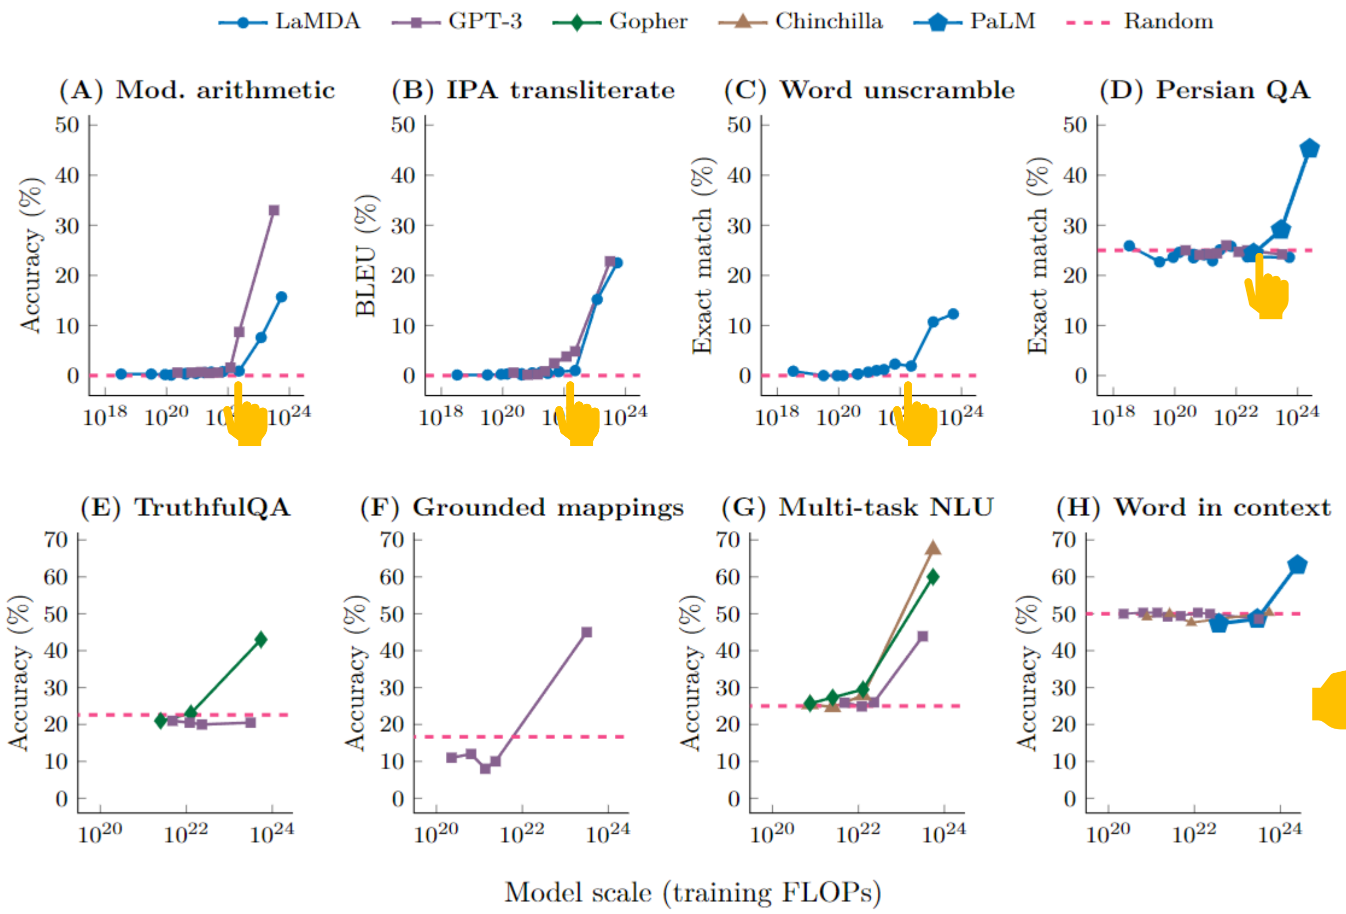

- **The easier reason: scaling law matters**
- **The harder version: emergent ability only emerges when the scale passes a certain scale threshold.**

Stronger out-of-distribution (O.O.D) task generalization emerges only after the scale passes a threshold.

Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., ... & Fedus, W. (2022). Emergent abilities of large language models. arXiv preprint arXiv:2206.07682.

The prevailing Bet in Embodied AI

➢Why O.O.D task generalization is so important in robotics?

- **Object variation** — shape, size, material, articulation, texture, mass
- **Scene variation** — layout, clutter, lighting, viewpoint, occlusion
- **Embodiment variation** — robot body, kinematics, controller, calibration
- **End-effector variation** — gripper, tool, contact geometry, compliance
- **Task variation** — goal, success criterion, task sequence, horizon
- **Instruction variation** — language phrasing, ambiguity, underspecification
- **Dynamics variation** — friction, deformability, motion, human interference

The prevailing Bet in Embodied AI

➢Why O.O.D task generalization is so important in robotics?

Even a slight change can turn a task collected in the training dataset into an OOD task.

EAI models, when at deployment, is inherently facing an open-world learning problem!

Leidner, Daniel Sebastian. Cognitive reasoning for compliant robot manipulation. Vol. 23. Heidelberg: Springer, 2019.

The prevailing Bet in Embodied AI

➢Have We Reached the Scaling Threshold in Embodied AI by large-scale pretraining?

- **Not yet**

Prof. Rodney Brooks

https://rodneybrooks.com/predictions-scorecard-2026-january-01/

The prevailing Bet in Embodied AI

➢Have We Reached the Scaling Threshold in Embodied AI by large-scale pretraining?

- **Not yet**

Prof. Rodney Brooks

Given that careful analysis from September I do not share the hype that surrounds humanoid robotics today. Some of it is downright delusional across many different levels.

To believe the promises of many CEOs of humanoid companies you have to accept the following conjunction.

1. Their robots have not demonstrated any practical work (I don't count dancing in a static environment doing exactly the same set of moves each time as practical work).
2. The demonstrated grasping, usually just a pinch grasp, in the videos they show is at a rate which is painfully slow and not something that will be useful in practice.
3. They claim that their robots will learn human-like dexterity but they have not shown any videos of multi-fingered dexterity where humans can and do grasp things that are unseen, and grasp and simultaneously manipulate multiple small objects with one hand. And no demonstrations of using the body with the hands which is how humans routinely carry many small things or one or two heavy things.

https://rodneybrooks.com/predictions-scorecard-2026-january-01/

The prevailing Bet in Embodied AI

➢Have We Reached the Scaling Threshold in Embodied AI by large-scale pretraining?

- **Not yet**

Prof. Rodney Brooks

11. The CEOs claim that there robots will be able to do everything, or many things, or a lot of things, that a human can do in just a few short years. They currently do none.
12. The CEOs claim a rate of adoption of these humanoid robots into homes and industries at a rate that is multiple orders of magnitude faster than any other technology in human history, including mainframe computers, and home computers and the mobile phones, and the internet. Many orders of magnitude faster. Here is a CEO of a humanoid robot company saying that they will be in 10% of US households by 2030. Absolutely no technology (even without the problems above) has ever come close to scaling at that rate.

https://rodneybrooks.com/predictions-scorecard-2026-january-01/

The prevailing Bet in Embodied AI

➢Have We Reached the Scaling Threshold in Embodied AI by large-scale pretraining?

- **Not yet**

Prof. Rodney Brooks

On December 25th the Wall Street Journal had a story headlined "Even the Companies Making Humanoid Robots Think They're Overhyped", with a lede of "Despite billions in investment, startups say their androids mostly aren't useful for industrial or domestic work yet". Here are the first two paragraphs of the story:

> Billions of dollars are flowing into humanoid robot startups, as investors bet that the industry will soon put humanlike machines in warehouses, factories and our living rooms.
>
> Many leaders of those companies would like to temper those expectations. For all the recent advances in the field, humanoid robots, they say, have been overhyped and face daunting technical challenges before they move from science experiments to a replacement for human workers.

And then they go on to quote various company leaders:

> "We've been trying to figure out how do we not just make a humanoid robot, but also make a humanoid robot that does useful work," said Pras Velagapudi, chief technology officer at Agility Robotics.

https://rodneybrooks.com/predictions-scorecard-2026-january-01/

The prevailing Bet in Embodied AI

➢In robotics, pretraining and post training are important, but NOT sufficient

Large-scale pretraining

keep scaling these will work!

Model

Data

Compute

Post-training & Adaptation

Domain specific SFT

RL

Reasoning

Long-horizon autonomy

- How far are we from the scaling threshold required for a GPT moment in robotics?
- **What if the scaling law for embodied AI is fundamentally different from that of NLP?**

The prevailing Bet in Embodied AI

➢**In robotics, pretraining and post training are important, but NOT sufficient**

- In robotics, pretraining and post-training are powerful, but they are not sufficient by themselves.
- Pretraining gives strong priors, but long-term embodied capability must continue to grow through interaction with the world.
- **The missing axis is experience.** Experience Scaling

The prevailing Bet in Embodied AI

➢In robotics, pretraining and post training are important, but NOT sufficient

- In robotics, pretraining and post-training are powerful, but they are not sufficient by themselves.
- Pretraining gives strong priors, but long-term embodied capability must continue to grow through interaction with the world.
- **The missing axis is experience.** 👉

THIS TALK TODAY:

(1) Two case studies from my own work

(2) How EAI Models Evolved & Future Trends

(3) A definition of Experience Scaling

(4) A draft of open research questions

Why O.O.D task generation in robotics is so hard?

Case study 1: diffusion adaptation and the pretraining ceiling

➢Test-time adaptation of pretrained diffusion models

- Pretrained model: a grasping task policy

⬇

- At test-time: generate collision-free trajectories without re-training.

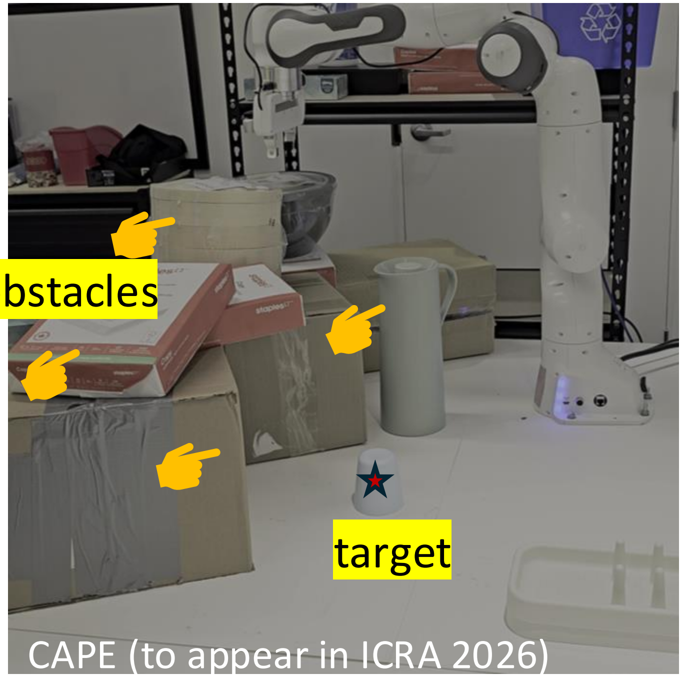

- **The idea (CAPE, to appear in ICRA 2026):**
  - Add context guidance in diffusion denoising when generating actions.

$$\tau_{t-1} = \mathrm{Denoise}_\theta(\tau_t, O) + \lambda \nabla_{\tau_t} L_{\mathrm{guid}}(\tau_t, O) + \sigma_t z$$

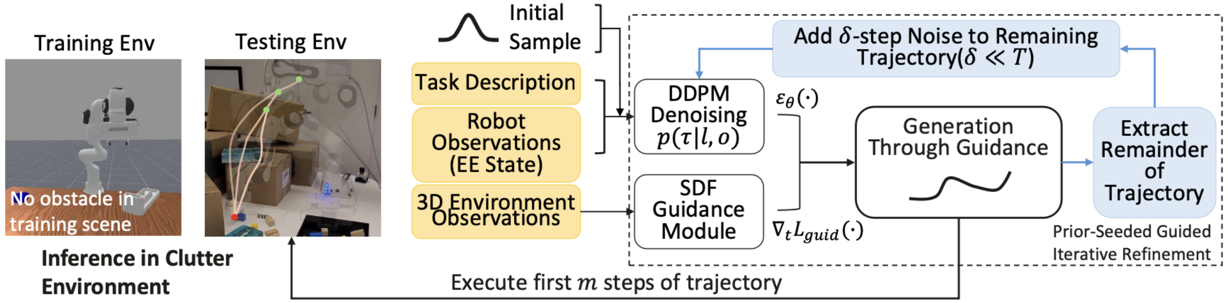

R. H. Yang, X. Zhao, L. M. Brunswic, M. Alban, M. Clemente, T. Cao, J. Jin, and A. Rasouli, "CAPE: Context-Aware Diffusion Policy Via Proximal Mode Expansion for Collision Avoidance," in IEEE International Conference on Robotics and Automation (ICRA), 2026.

Case study 1: diffusion adaptation and the pretraining ceiling

➢Test-time adaptation of pretrained diffusion models

- Pretrained model: a grasping task policy

⇩

- At test-time: generate collision-free trajectories without re-training.

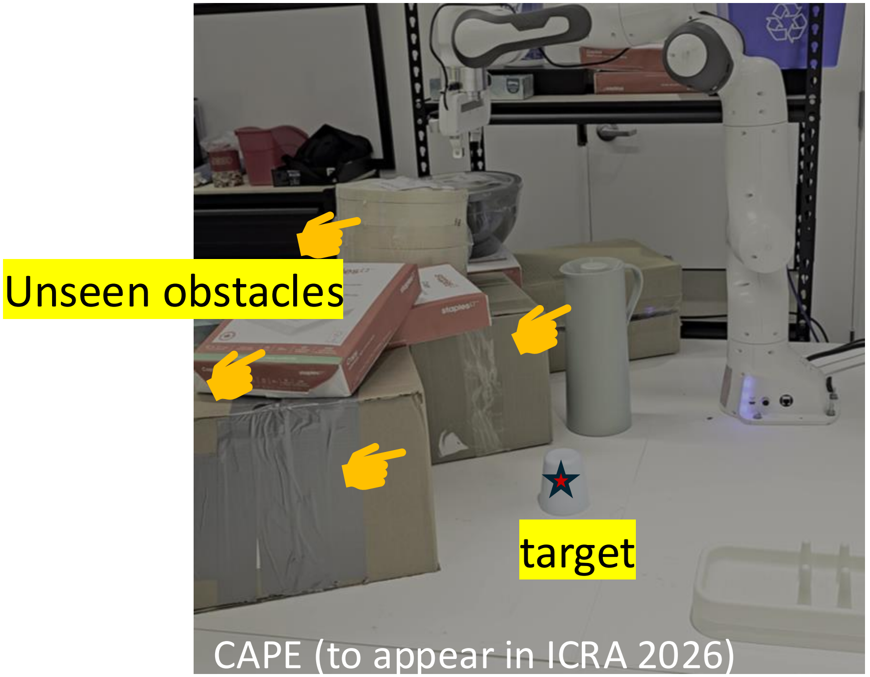

- **The idea (CAPE, to appear in ICRA 2026):**
  - Add collision-free gradient guidance in diffusion denoising steps when generating actions.

$$\tau_{t-1} = \mathrm{Denoise}_\theta(\tau_t, O) + \lambda \nabla_{\tau_t} L_{\mathrm{guid}}(\tau_t, O) + \sigma_t z$$

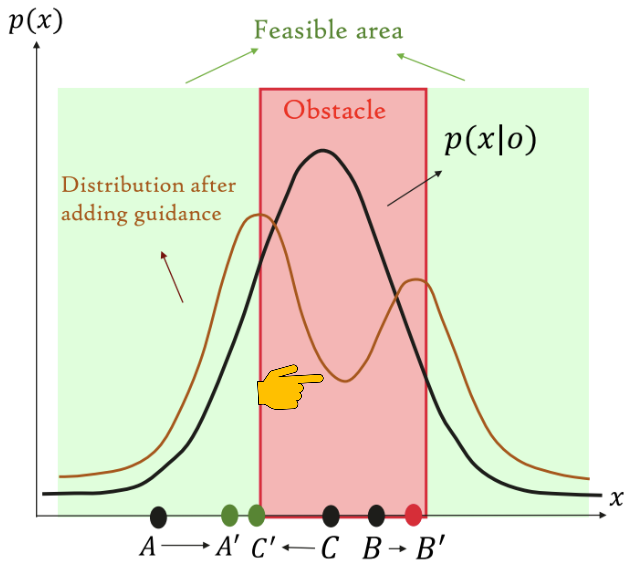

R. H. Yang, X. Zhao, L. M. Brunswic, M. Alban, M. Clemente, T. Cao, J. Jin, and A. Rasouli, "CAPE: Context-Aware Diffusion Policy Via Proximal Mode Expansion for Collision Avoidance," in IEEE International Conference on Robotics and Automation (ICRA), 2026.

Case study 1: diffusion adaptation and the pretraining ceiling

➢Test-time adaptation of pretrained diffusion models

- Too weak guidance fails to move the trajectory enough.
- Too strong guidance distorts the sample off-distribution.
- **CAPE preserve a structured prior while interactively expanding context-aware mode support.**

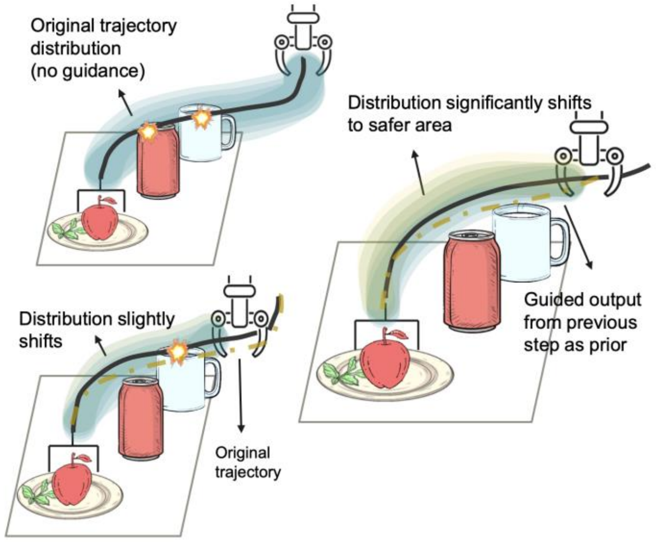

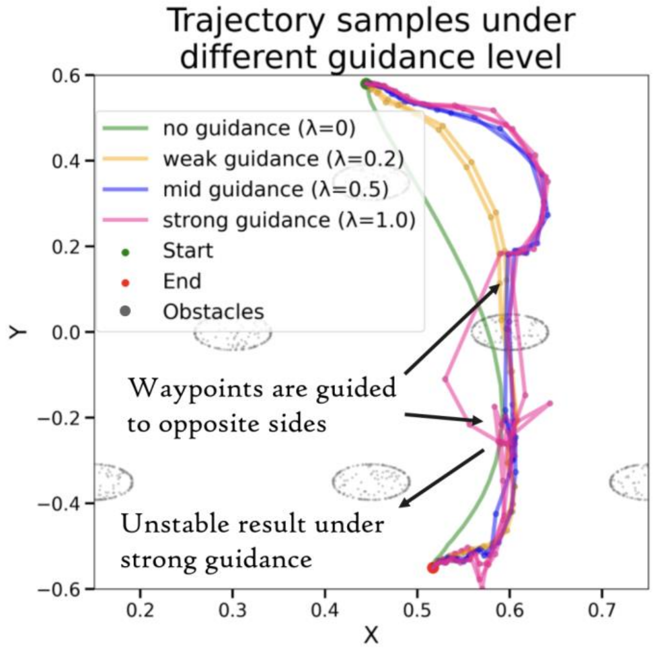

LIMITATIONS

Test-time adaptation is still bounded by the support of the pretrained action distribution.

- Although diffusion models can adapt at test time, their adaptation remains constrained by the support of the learned action distribution.
- Guidance can shift the trajectory, but it cannot create missing action modes from nothing.

R. H. Yang, X. Zhao, L. M. Brunswic, M. Alban, M. Clemente, T. Cao, J. Jin, and A. Rasouli, "CAPE: Context-Aware Diffusion Policy Via Proximal Mode Expansion for Collision Avoidance," in IEEE International Conference on Robotics and Automation (ICRA), 2026.

Case study 2: cross-domain transfer and the pretraining ceiling

➢Transfer from Human ADL data to Surgical Robots

- Surgical data: scares, safety and ethical constraints.
- Human Activities of Daily Living task data: easy to obtain, much diverse motor skills.

- **The idea:**
  - Modular Deep Successor Features encode cross-embodiment reusable motor structures

Human daily dataset [7]

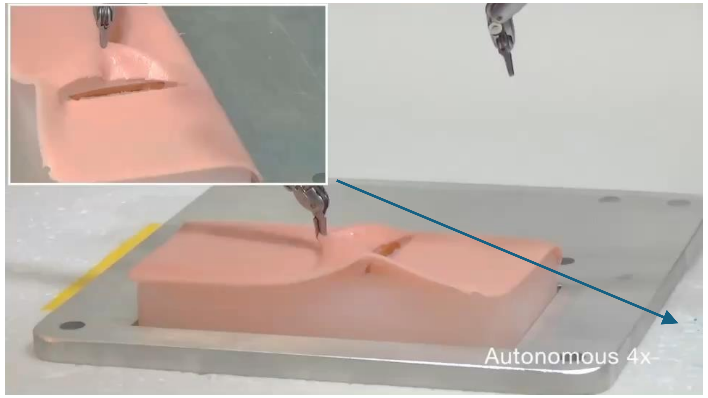

Robotic surgical task [23]

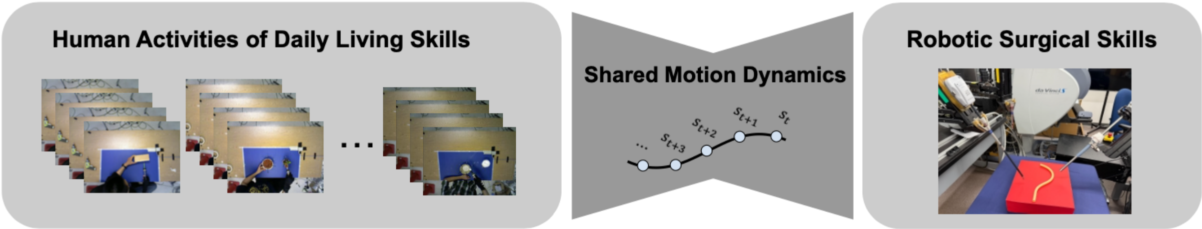

Cross-embodiment transfer problem

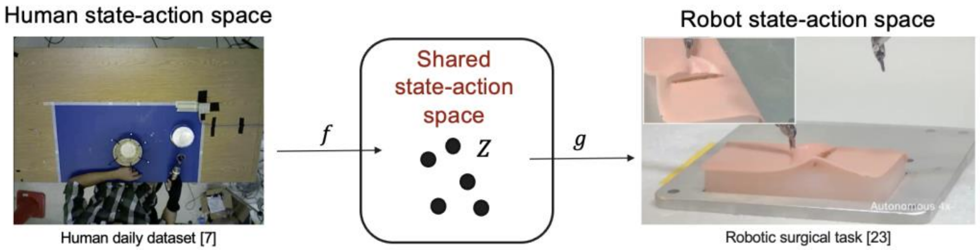

Cross-task transfer problem

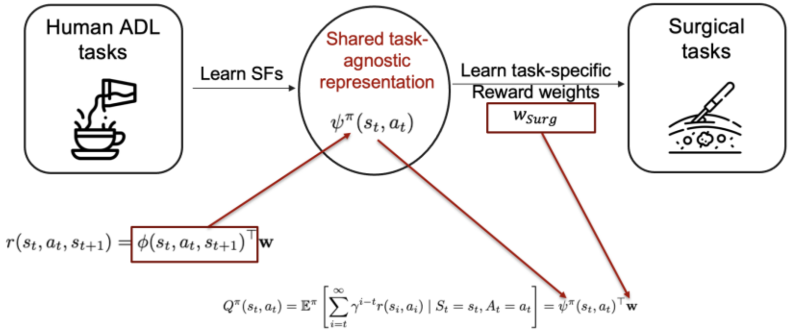

Y. Hu, M. Tavakoli, and J. Jin, "Pretraining using comparable human activities of daily living dataset in robotic surgical task learning," IEEE Transactions on Medical Robotics and Bionics, 2025.

Case study 2: cross-domain transfer and the pretraining ceiling

➢Transfer from Human ADL data to Surgical Robots

- Naïve direct transfer can fail under embodiment and task mismatch
- Our modular SF learns cross-task/cross-embodiment motor structures, thus transfer works much better.
- Our modular SF can automatically select skills by learning the weights when transferred in the surgical tasks.

Modular Successor Feature Framework

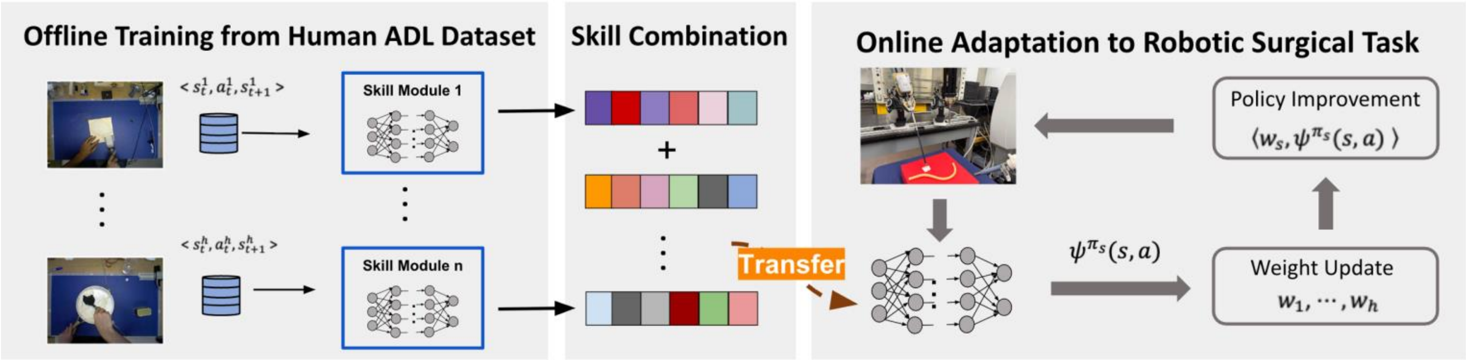

LIMITATIONS

Cross-domain transfer is still bounded by shared sensorimotor structure.

- Structured representation help domain and task transfer.
- But it can only help if the source and target domain share the right transition structure.
- ADL tasks with minimal relevance to the surgical domain yield little or no transfer benefit.

Y. Hu, M. Tavakoli, and J. Jin, “Pretraining using comparable human activities of daily living dataset in robotic surgical task learning,” IEEE Transactions on Medical Robotics and Bionics, 2025.

These are all small-scale models.

What if we scale up model size, dataset size, and compute throughput?

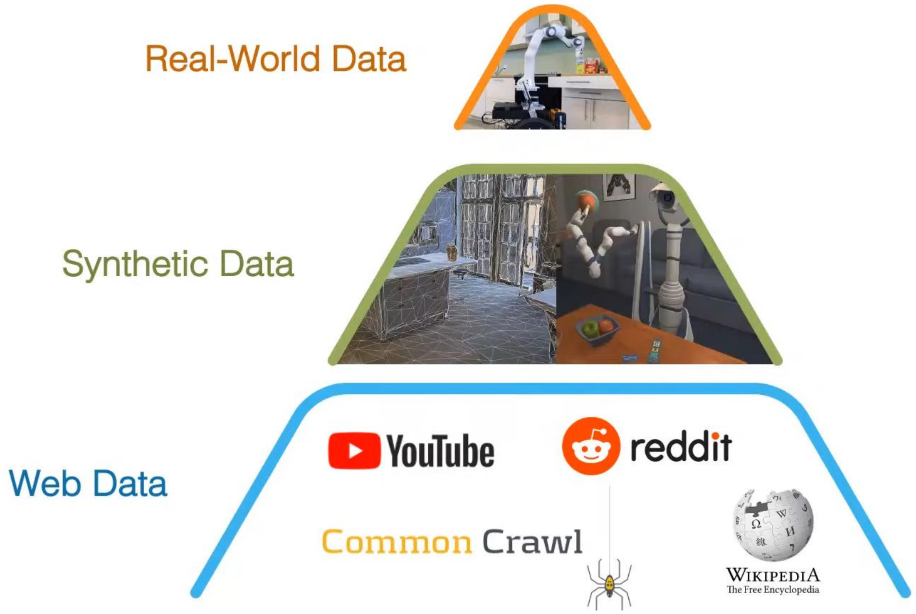

The Data Pyramid Zhu et al. 2022

Lin, Fanqi, et al. "Data scaling laws in imitation learning for robotic manipulation." *arXiv preprint arXiv:2410.18647* (2024).

How the EAI Models Evolved since 2022

➢Phase 1: Fully End-to-End Generalist Ambition

(1) 2022: Small transformer-block based models

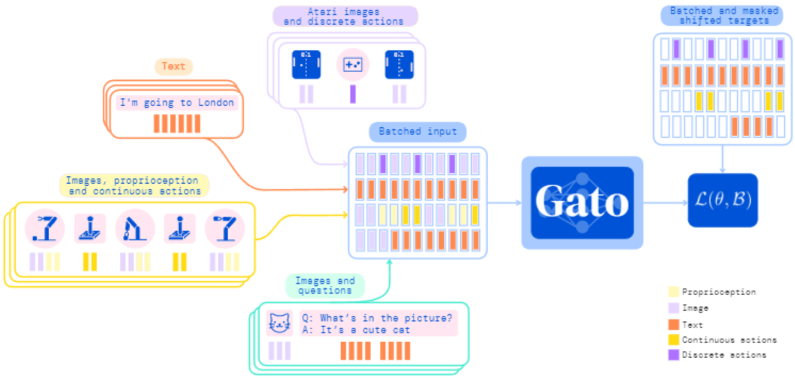

Gato, Reed et al. 2022

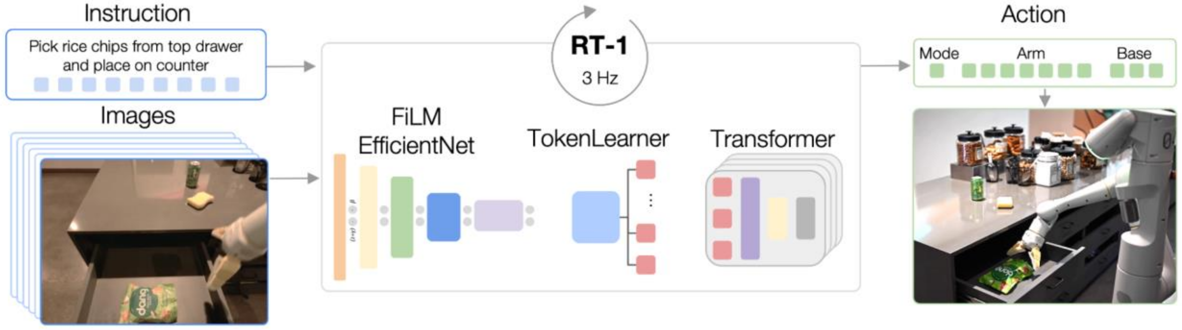

RT1, Brohan et al. 2022

~ 300M to 1.2B model params size

~ 35M model params size

Reed, Scott, et al. "A generalist agent." *arXiv preprint arXiv:2205.06175* (2022).
Brohan, Anthony, et al. "Rt-1: Robotics transformer for real-world control at scale." *arXiv preprint arXiv:2212.06817* (2022).

How the EAI Models Evolved since 2022

➢Phase 1: Fully End-to-End Generalist Ambition

(2) 2023: Large VLM based models

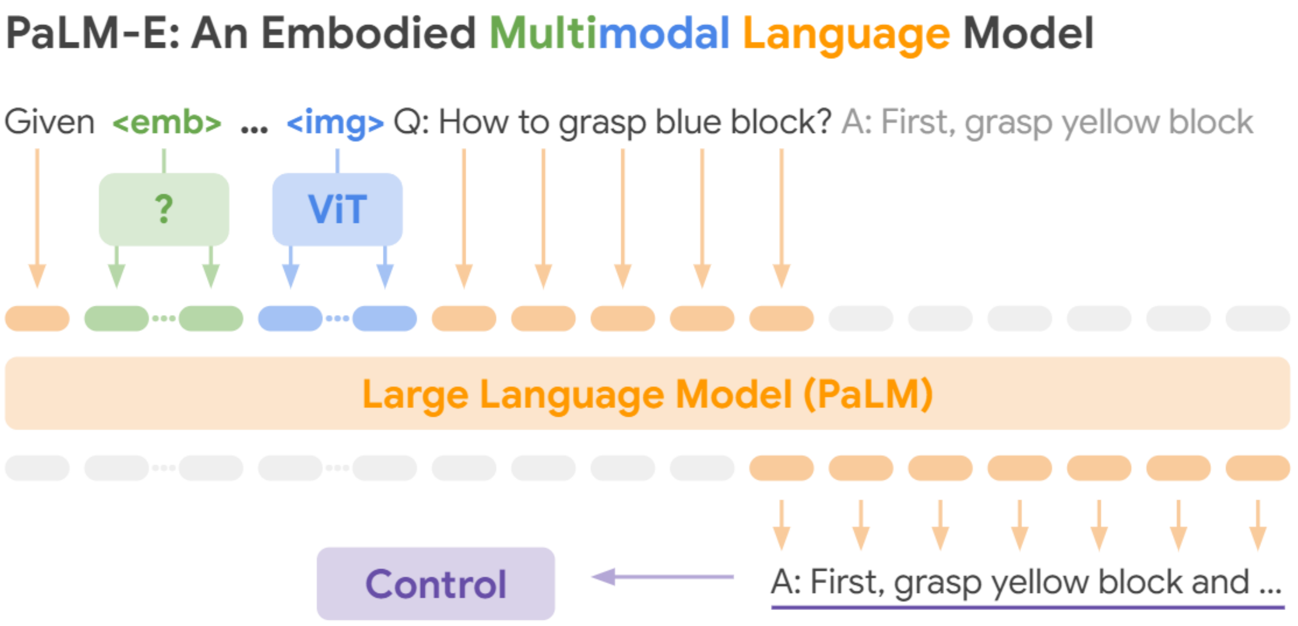

PaLM-E, Driess et al. 2023

~ 562B model params size

Driess, Danny, et al. "Palm-e: An embodied multimodal language model." (2023).

How the EAI Models Evolved since 2022

➢Lessons from Phase I

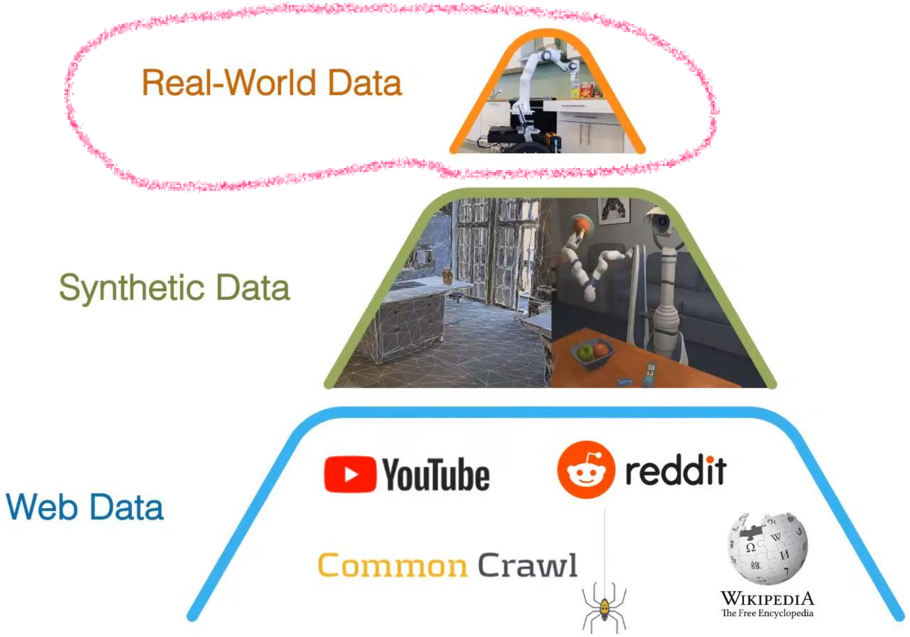

The Data Pyramid Zhu et al. 2022

Lessons learned from phase 1

Mixed end2end training helps embodied reasoning, but real-world robot data matters more.

- Fully mixed end-to-end pretraining did not automatically yield strong embodied competence.
- The benefit of scale depended strongly on data relevance to robot control.
- General multimodal priors helped, but did not remove embodiment and execution constraints.
- The field learned that scaling in robotics is also a data-quality problem.

Reed et al., 2022; Brohan et al., 2022; Driess et al., 2023

Reed, Scott, et al. "A generalist agent." *arXiv preprint arXiv:2205.06175* (2022).
Brohan, Anthony, et al. "Rt-1: Robotics transformer for real-world control at scale." *arXiv preprint arXiv:2212.06817* (2022).
Driess, Danny, et al. "Palm-e: An embodied multimodal language model." (2023).

How the EAI Models Evolved since 2022

➢Phase 2: 0.5 End-to-End Models / Vision-Language-Action (VLA) models

The Rationale behind the VLAs:

- If real-world robot data drives the strongest gains, why not keep the pretrained VLM backbone fixed and train a robot-specific action expert?
- This motivates a modular VLA design: retain the pretrained VLM backbone, and train a robot-data-specialized action expert for control.

Vision-Language-Action (VLA) models ⟺ **Pretrained VLM + Action Expert (diffusion/transformer)**

How the EAI Models Evolved since 2022

➢Phase 2: 0.5 End-to-End Models / Vision-Language-Action (VLA) models

(1) Transformer-based Action Expert

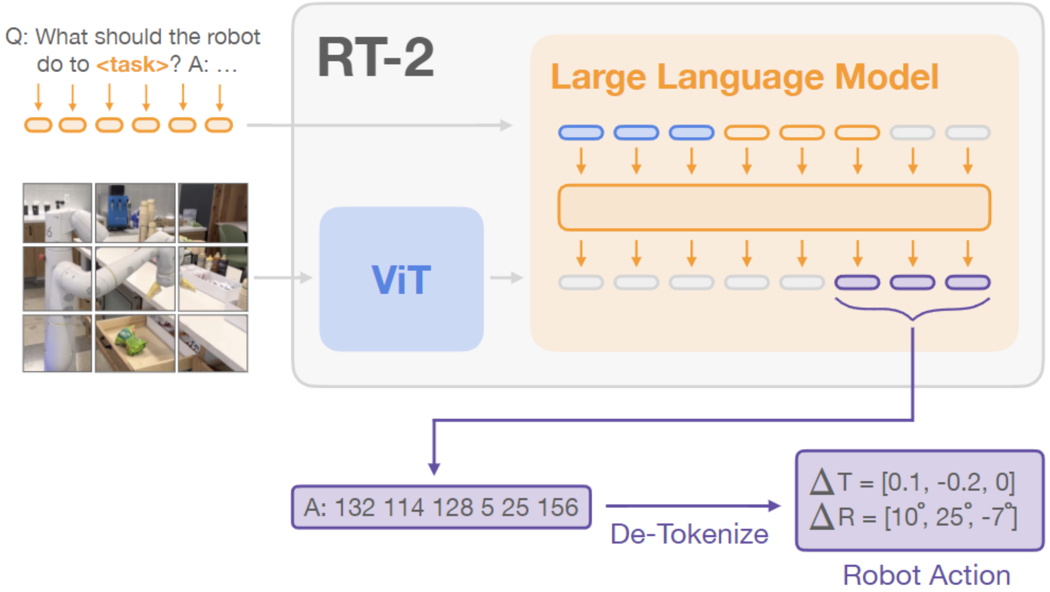

RT2, Brohan et al. 2023

- PaLM-E VLM ~ 12B model params size
- PaLI-X VLM ~ 55B model params size

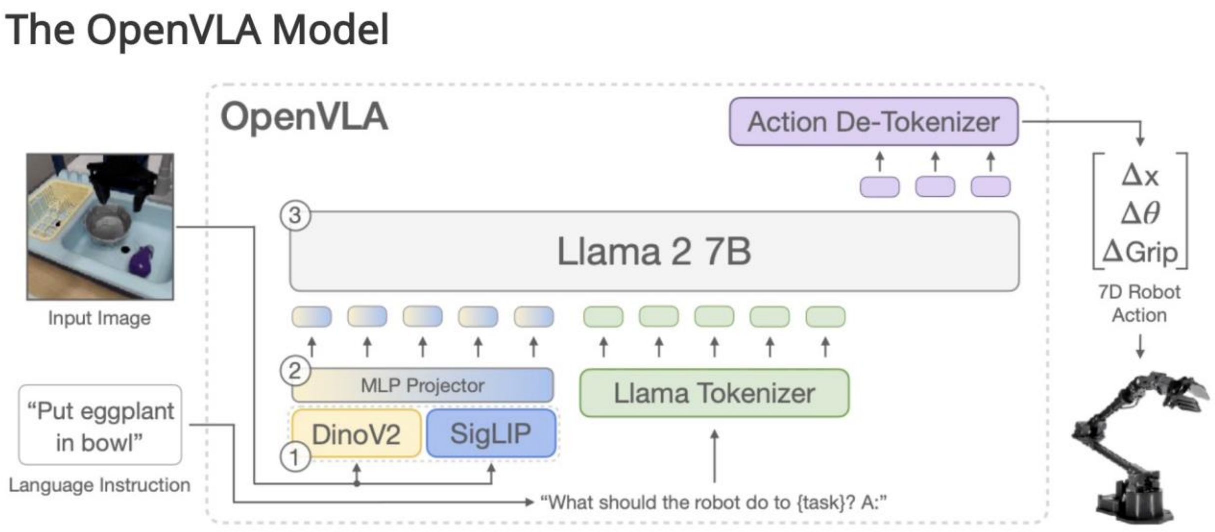

OpenVLA, Kim et al. 2024

~ 7B params size

Brohan, Anthony, et al. "Rt-2: Vision-language-action models transfer web knowledge to robotic control." *arXiv preprint arXiv:2307.15818* (2023).
Kim, Moo Jin, et al. "Openvla: An open-source vision-language-action model." *arXiv preprint arXiv:2406.09246* (2024).

How the EAI Models Evolved since 2022

➢Phase 2: 0.5 End-to-End Models / Vision-Language-Action (VLA) models

(1) Transformer-based Action Expert

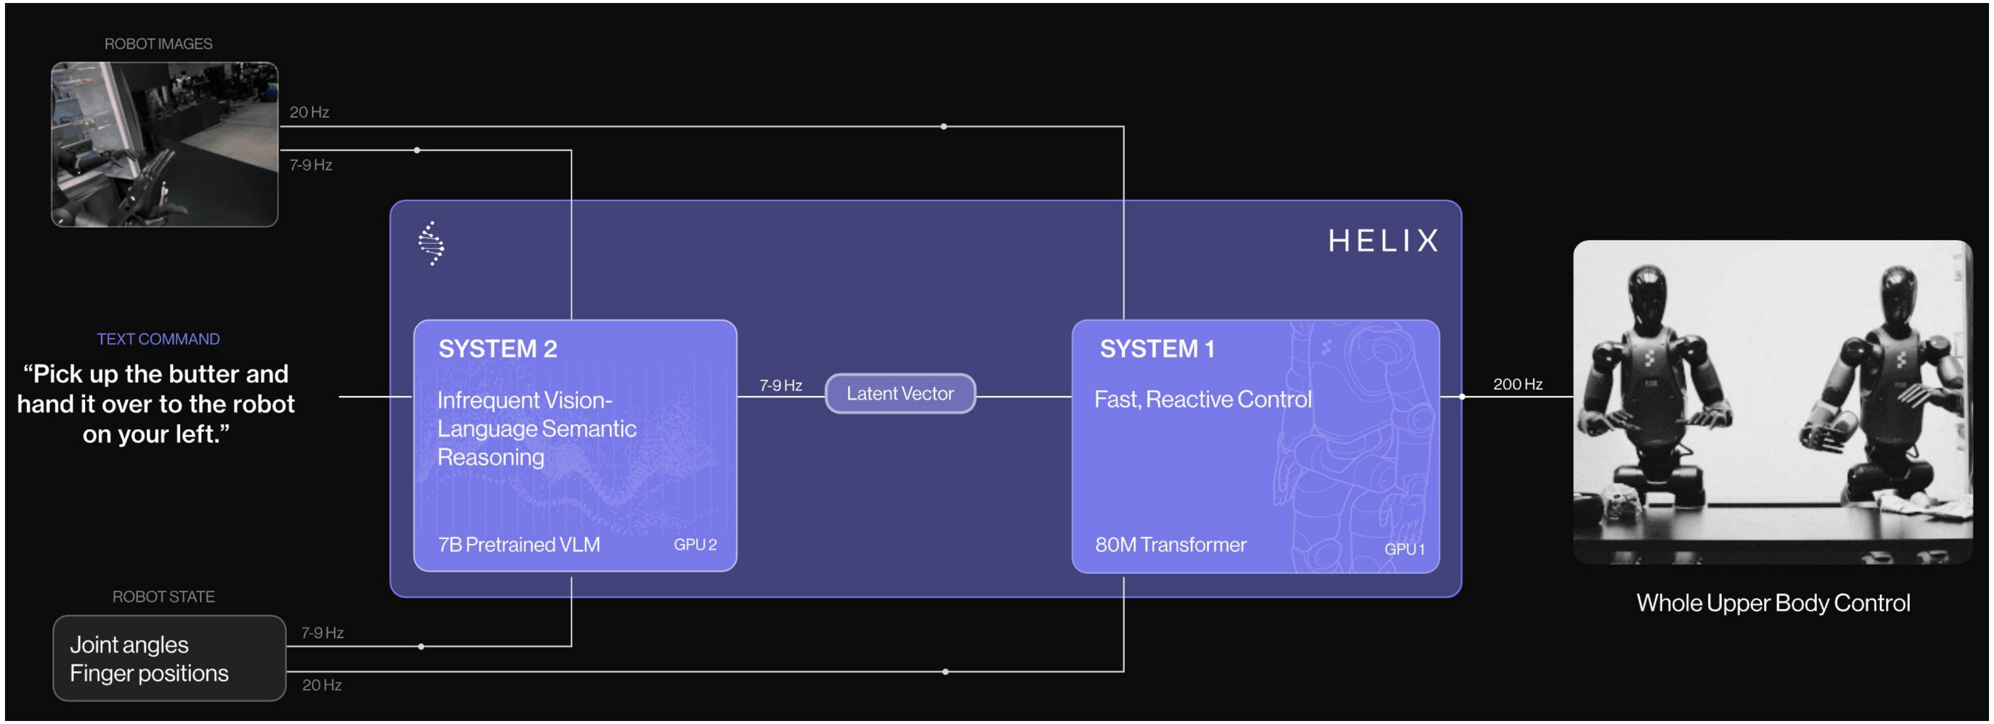

https://www.figure.ai/news/helix

How the EAI Models Evolved since 2022

➢Phase 2: 0.5 End-to-End Models / Vision-Language-Action (VLA) models

(2) Diffusion-based Action Expert ➡ much better high DOF / continuous control

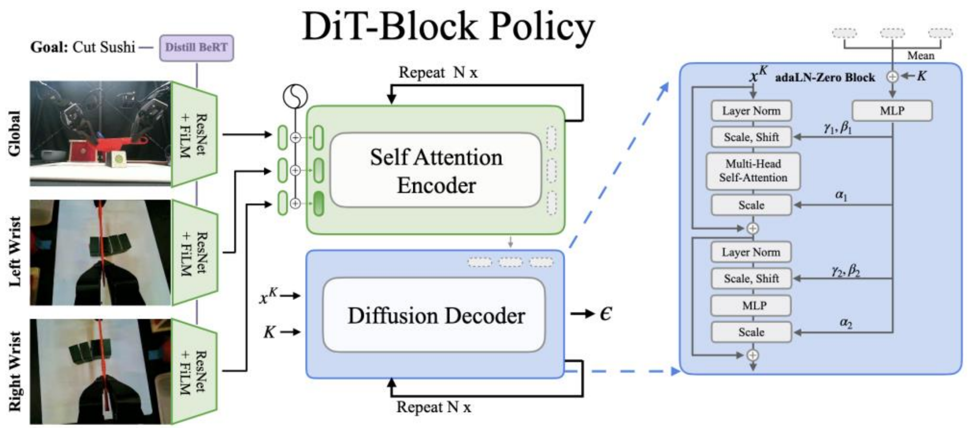

DiT, Dasari et al. 2024

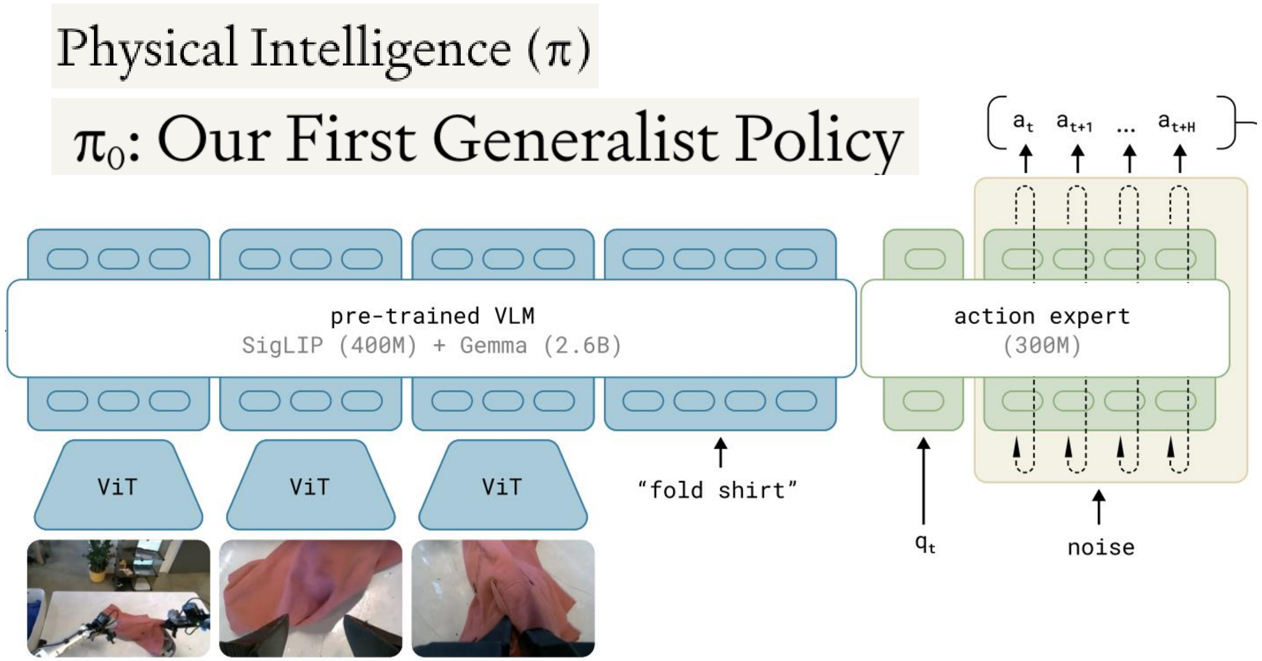

Pi-0, Black et al. 2024

~ 3B params size

Dasari, Sudeep, et al. "The ingredients for robotic diffusion transformers." *arXiv preprint arXiv:2410.10088* (2024).

Black, Kevin, et al. "π0: A vision-language-action flow model for general robot control, 2024." *URL https://arxiv. org/abs/2410.24164.*

Pretraining Dataset composition of Pi-zero model

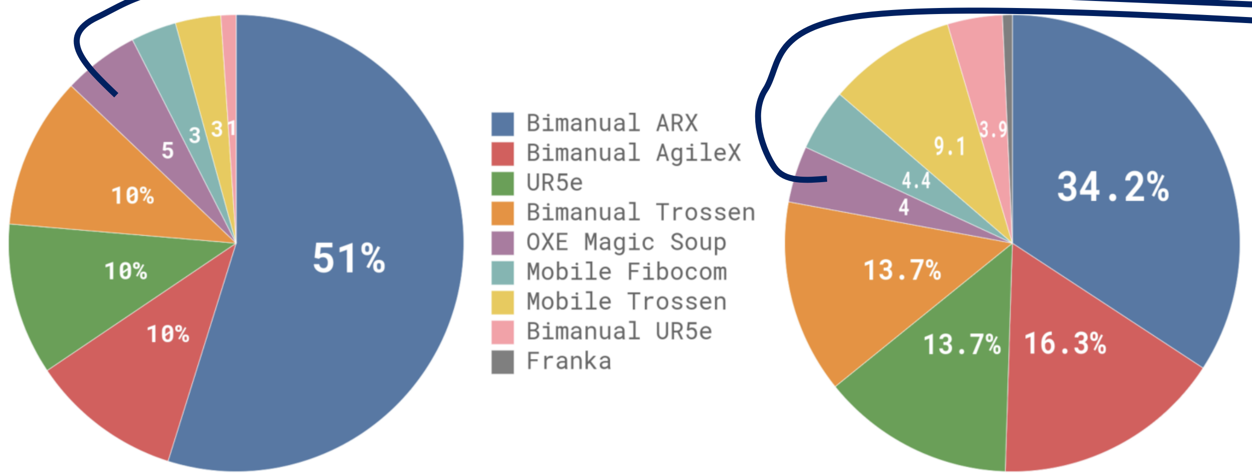

Fig. 4: **Overview of our dataset**: The pre-training mixture consists of a subset of OXE [10] and the π dataset. We use a subset of OXE, which we refer to as OXE Magic Soup [24]. The left figures illustrates the weight of the different datasets in the pre-training mixture. The right figure illustrates their relative sizes as measured by the number of steps.

Pi-0, Black et al. 2024

OpenVLA's Magic Soup
970k robot demonstration trajectories
~ 5% of the dataset

Overall dataset in pre-training:
19.4M trajectories

~95% real-world robot dataset collected by the startup company themselves. Why?

Black, Kevin, et al. "π0: A vision-language-action flow model for general robot control, 2024." *URL https://arxiv. org/abs/2410.24164.*

Pretraining Dataset composition of Pi-zero model

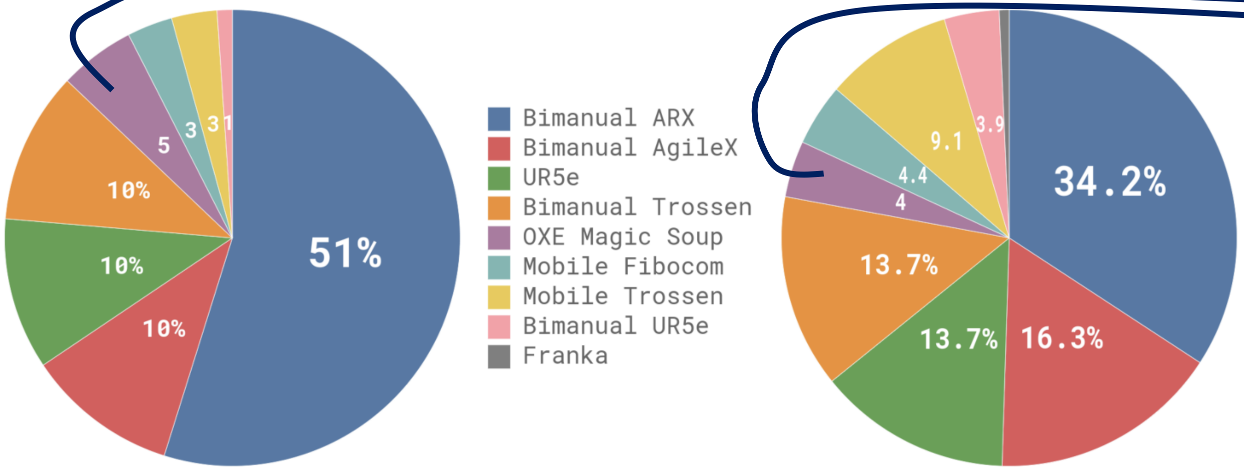

Fig. 4: **Overview of our dataset**: The pre-training mixture consists of a subset of OXE [10] and the π dataset. We use a subset of OXE, which we refer to as OXE Magic Soup [24]. The left figures illustrates the weight of the different datasets in the pre-training mixture. The right figure illustrates their relative sizes as measured by the number of steps.

Pi-0, Black et al. 2024

OpenVLA Magic Soup
970k robot demonstration trajectories
~ 5% of the dataset

Overall dataset in pre-training:
19.4M trajectories

The field found out that when training VLA models:

1. Cross-domain dataset seems not really help.
2. Simulation data not really help either.
3. Only real-world robot data improves the model significantly!

Black, Kevin, et al. "π0: A vision-language-action flow model for general robot control, 2024." *URL https://arxiv. org/abs/2410.24164*.

How the EAI Models Evolved since 2022

➢Lessons from Phase 2

Still, there is no clear evidence that we have reached the scaling threshold required for strong OOD task generalization.

Lessons learned from phase 2

High-quality robot data are still the critical scarce resource. EAI models are SUPER data hungry!

- The VLM inside a VLA is not fully utilized for reasoning and language understanding.
  - Using a latent vector as the conditional variable for action generation shows bottlenecks that it sacrifices part of the earlier full end-to-end expressivity.
- High-quality real-world robot data is the first choice; simulator data were useful, but still secondary.
- Realize that VLAs are just strong backbones, domain-specific fine-tuning (SFT/RL) is necessary.

How the EAI Models Evolved since 2022

➢Phase 3: Seeking Beyond Robot Data & The Return of End-to-end Models

- **We still do not know how large a robot dataset must be to reach the scaling threshold.**
- **The long tail of real-world interaction is unbounded.**

token scale in LLM ≠ token scale in EAI

- **Object variation** — shape, size, material, articulation, texture, mass
- **Scene variation** — layout, clutter, lighting, viewpoint, occlusion
- **Embodiment variation** — robot body, kinematics, controller, calibration
- **End-effector variation** — gripper, tool, contact geometry, compliance
- **Task variation** — goal, success criterion, task sequence, horizon
- **Instruction variation** — language phrasing, ambiguity, under specification
- **Dynamics variation** — friction, deformability, motion, human interference

Leidner, Daniel Sebastian. Cognitive reasoning for compliant robot manipulation. Vol. 23. Heidelberg: Springer, 2019.

How the EAI Models Evolved since 2022

➢Phase 3: Seeking Beyond Robot Data & The Return of End-to-end Models

- **High-quality real robot data are too scare relative to the ambition of general-purpose robots.**

And EAI data collection is costly!

Local governments in Shanghai have funded 40 training centers to address a data shortage in robotics research.

A simpler but unstable data collection solution, UMI (Chi et al. 2024)

https://restofworld.org/2026/china-robots-training-centers-workers/

https://arxiv.org/abs/2402.10329

How the EAI Models Evolved since 2022

➢Phase 3: Seeking Beyond Robot Data & The Return of End-to-end Models

- **Phase 1 proves that all mixture of the data pyramid won't work.**
- **Phase 2 realizes that only high-quality real-world robot data matter the most, but are scare.**

Beyond Robot data

where can more usable embodied data come from?

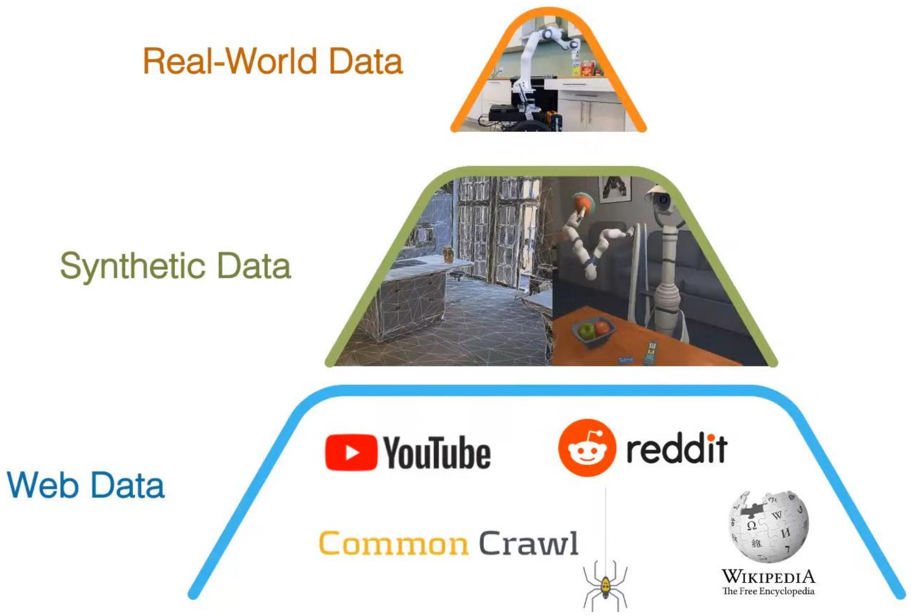

The Data Pyramid Zhu et al. 2022

Any data that are robot-relevant

- Web data that related to robotic tasks
- Image/video QA that support robotic reasoning/spatial understanding.
- Object detection / keypoint prediction / 2D trajectories
- Multi-view correspondence
- Success detection / Instrument reading
- Human Activity Videos
- Synthetic Data

Beyond Robot data

where can more usable embodied data come from?

Any data that are robot-relevant

- Web data that related to robotic tasks
- Image/video QA that support robotic reasoning/spatial understanding.
- Object detection / keypoint prediction / 2D trajectories
- Multi-view correspondence
- Success detection / Instrument reading
- **Human Activity Videos**
- **Synthetic Data**

π-0.5 (Black, et al. 2025)

Gemini Robotics-ER, 1.5&1.6 (Abeyruwan et al. 2025, Graesser et al. 2026)

GR00T N1 (Bjorck et al. 2025)

World Models (Ali et al. 2025, …)

World Action Models (Wang et al. 2025)

…

Beyond Robot data

where can more usable embodied data come from?

Any data that are robot-relevant

- Web data that related to robotic tasks
- Image/video QA that support robotic reasoning/spatial understanding
- Object detection / keypoint prediction / 2D trajectories
- Multi-view correspo
- Success detection / Instrument reading
- Human Activity Videos
- Synthetic Data

π-0.5 (Black, et al. 2025)

eyruwan et al. 2025, Graesser et al. 2026)

2025)

2025, ...)

t al. 2025)

...

wait...

how to train them in one unified model together with robot trajectories?

VLM + Task Orchestration,
Inverse Dynamics Modeling (IDM),
Latent Action Models (LAM),
...

Sorry I won't cover the model details here due to time limits.

➤ **Phase 3: Seeking Beyond Robot Data & The Return of End-to-end Models**

Beyond Robot data

where can more usable embodied data come from?

Any data that are robot-relevant

- Web data that related to robotic tasks
- Image/video QA that support robotic reasoning/spatial understanding.
- Object detection / keypoint prediction / 2D trajectories
- Multi-view correspondence
- Success detection / Instrument reading
- Human Activity Videos
- Synthetic Data

π-0.5 (Black, et al. 2025)

Gemini Robotics-ER, 1.5&1.6 (Abeyruwan et al. 2025, Graesser et al. 2026)

GR00T N1 (Bjorck et al. 2025)

World Models (Ali et al. 2025, …)

World Action Models (Wang et al. 2025)

…

What's next BIG thing in phase 3?

➢ **Phase 3: Seeking Beyond Robot Data & The Return of End-to-end Models**

Beyond Robot data

where can more usable embodied data come from?

Any data that are robot-relevant

- Web data that related to robotic tasks
- Image/video QA that support robotic reasoning/spatial understanding.
- Object detection / keypoint prediction / 2D trajectories
- Multi-view correspondence
- Success detection / Instrument reading
- Human Activity Videos
- Synthetic Data

More curated robot-relevant data:

- ❑ Geometry Supervisions
- ❑ Affordance
- ❑ Occlusion/ Scene Depth Predictions
- ❑ Tactile Signal Activate Predictions
- ❑ Contact / Collision Predictions
- ❑ …

π-0.5 (Black, et al. 2025)

Gemini Robotics-ER, 1.5&1.6 (Abeyruwan et al. 2025, Graesser et al. 2026)

GR00T N1 (Bjorck et al. 2025)

World Models (Ali et al. 2025, …)

World Action Models (Wang et al. 2025)

…
Model
?

The Next BIG thing in EAI Phase 3

How the EAI Models Evolved since 2022

➢ **The hidden line through all of these: chasing scaling law through EAI data**

> **The history of embodied AI since 2022 is the history of chasing a scaling regime for robotics, and therefore repeatedly redesigning both models and data pipelines in search of more and more EAI data.**

scaling threshold

Initial attempt

PaLM-E, Driess et al. 2023
RT1, Brohan et al. 2022
Gato, Reed et al. 2022

Phase 1: 2022, 2023

High-quality robot data matters!

Pi-0, Black et al. 2024
Figure AI, Helix model
OpenVLA, Kim et al. 2024
RT2, Brohan et al. 2023

Phase 2: 2023~now

All robot-relevant data are useful.

π-0.5 (Black, et al. 2025)
Gemini Robotics-ER, 1.5&1.6 (Abeyruwan et al. 2025, Graesser et al. 2026)
GR00T N1 (Bjorck et al. 2025)
World Models (Ali et al. 2025, …)
World Action Models (Wang et al. 2025)

Phase 3: 2025~now

How large of the data/model/compute is considered “large” enough to cross the threshold of the scaling law in these EAI mo

IEEE Spectrum FOR THE TECHNOLOGY INSIDER

GUEST ARTICLE | ROBOTICS

28 May 2024

Will Scaling Solve Robotics?

- If yes, follow my design in Phase3, and we will solve robotics.
- End of the game!

How the EAI Models Evolved since 2022

➢The “elephant in the room”: the scaling law in robotics is less settled

The “scaling law” paper, Lin et al. 2024: ICLR 2025 Oral, Best paper in workshop CoRL 2024

Can we trust in THIS “scaling law” for robotics with only a limited task configurations?

Data Scaling Laws in Imitation Learning for Robotic Manipulation

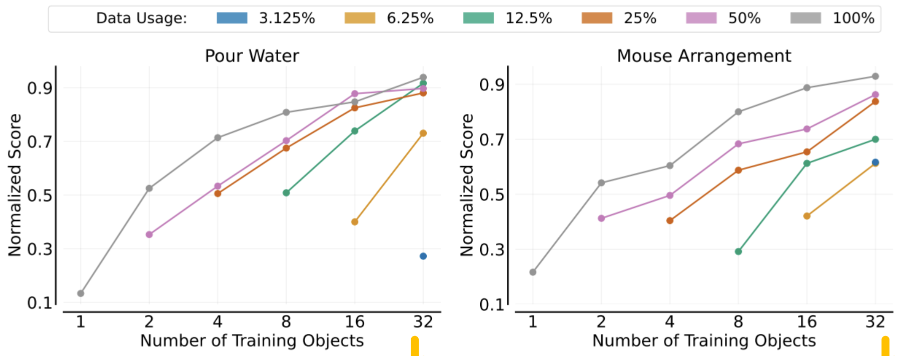

Figure 2: **Object generalization.** Each curve corresponds to a different fraction of demonstrations used, with normalized scores shown as a function of the number of training objects.

VS

Our everyday activities are much more diverse!

Grauman, Kristen, et al. "Ego4d: Around the world in 3,000 hours of egocentric video." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2022.

How the EAI Models Evolved since 2022

➢Phase 3: Seeking Beyond Robot Data & The Return of End-to-end Models

- **We still do not know how large a robot dataset must be to reach the scaling threshold.**
- **The long tail of real-world interaction is effectively unbounded.**

The field has not yet shown that static pretraining alone can deliver open-world embodied intelligence.

- **Object variation** — shape, size, material, articulation, texture, mass
- **Scene variation** — layout, clutter, lighting, viewpoint, occlusion
- **Embodiment variation** — robot body, kinematics, controller, calibration
- **End-effector variation** — gripper, tool, contact geometry, compliance
- **Task variation** — goal, success criterion, task sequence, horizon
- **Instruction variation** — language phrasing, ambiguity, under specification
- **Dynamics variation** — friction, deformability, motion, human interference

Leidner, Daniel Sebastian. Cognitive reasoning for compliant robot manipulation. Vol. 23. Heidelberg: Springer, 2019.

A Different Path: From pretraining to experience-scaling

➢If static pretraining is NOT sufficient, what are alternative pathways?

- The problem may be deeper than chasing a sufficient static scale of data, model size & compute.
- In embodied settings, the agent is bounded, but the world is effectively unbounded.
- Robotic interaction is multimodal, contact-rich, path-dependent, and long-horizon.
  - New objects, layouts, materials, failure modes, and human preferences continually create new OOD conditions.
- **Data, models, and compute alone may not be the whole answer.**

THE CANADIAN PRESS/Nathan Denette

A Different Path: From pretraining to experience-scaling

➢ **If static pretraining is NOT sufficient, what are alternative pathways?**

The Big World Hypothesis (BWH), Javed et al. 2024

- The big world hypothesis says that for many learning problems, the world is multiple orders of magnitude larger than the agent.
- The agent neither fully perceives the state of the world nor can it learn the correct value or optimal action for each state.
- It has to rely on approximate solutions to achieve its goals.

THE CANADIAN PRESS/Nathan Denette

A Different Path: From pretraining to experience-scaling

➢If static pretraining is NOT sufficient, what are alternative pathways?

- **Will post-training (SFT & RL) be the rescue?**
  - SFT and RL post-training are not a satisfactory general solution either, because they are difficult to scale across continually changing OOD shifts.
- Can we think about the problem in this way?
  - Pretraining gives strong priors, we agree on that.
  - But what is missing is a way for the robot’s embodied capability to keep growing through interaction with the unbounded world under bounded resources (compute, energy consumption).
  - **The missing axis here is experience!**

image source: dreamstime.com

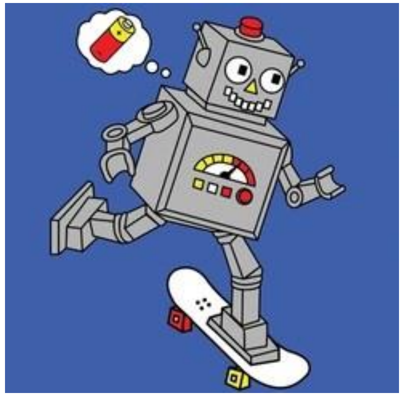

A Different Path: From pretraining to experience-scaling

➢What Counts as Experience?

- **Experience = Closed-Loop transition tuples when the robot interacts with the world**

next observation

observation action reward/outcome feedback

$$e_t = (o_t, a_t, o_{t+1}, r_t, h_t, c_t, \Delta_t)$$

human correction or instruction

resource cost (energy, compute)

novelty, or learning signal from pretrained priors

A Different Path: From pretraining to experience-scaling

➢Effective Experience: Not All Experience is Useful

- **Experience Scaling should scale with only the amount of effective embodied experience.**
  - 10,000 repeated easy successes may contribute very little effective experience.
  - 1 failure followed by human instruction, external EAI model correction, and recovery may contribute much more.

If the robot learns how to reach a target, repeating the same reaching traj 1k times won't help future learning.

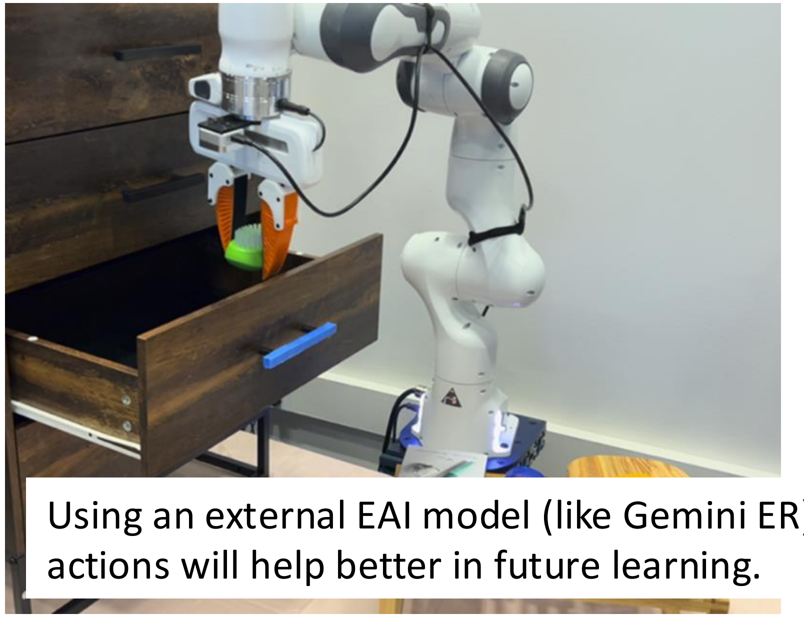

Using an external EAI model (like Gemini ER) to correct colliding actions will help better in future learning.

https://www.futurity.org/manipgen-robots-manipulating-objects-3259572-2/

https://www.cs.cmu.edu/news/2024/manipgen

A Different Path: From pretraining to experience-scaling

➢Effective Experience: Not All Experience is Useful

- **Effective Experience:**
  - Measures how much closed-loop interaction contributes to future reusable skills, adaptation, and regret reduction, rather than how many transitions the agent has stepped.
- **A function measures the value of an experience e_t :**

$$w(e_t)$$

High-value interaction: $w(e_t) > 0$

- novelty / informative representations
- bottleneck states
- subproblem discovery
- skill refinement
- failure recovery skills
- human preference update
- prediction correction

Human: throw away the snack on the close counter
Robot: 1. go to the close counter, and ask:
What snacks are on the counter?
Human: actually I changed my mind, I want you to throw away something on the table
2. go to the table and ask:
What snacks are on the table?
Human: nevermind i want you to finish your previous task
3. go to the close counter and ask:
What snacks are on the counter?
human: apple and chip bag
4. pick up the apple and continue
5. go to the trash and continue

https://innermonologue.github.io/

A Different Path: From pretraining to experience-scaling

➢Effective Experience: Not All Experience is Useful

- **Effective Experience:**
  - Measures how much closed-loop interaction contributes to future reusable skills, adaptation, and regret reduction, rather than how many transitions the agent has stepped.
- **A function measures the value of an experience e_t :**

$$w(e_t)$$

low-value interaction: $w(e_t) \approx 0$

- non-informative steps
- e.g., redundant execution with no change in policy, representation, skill library, or memory.

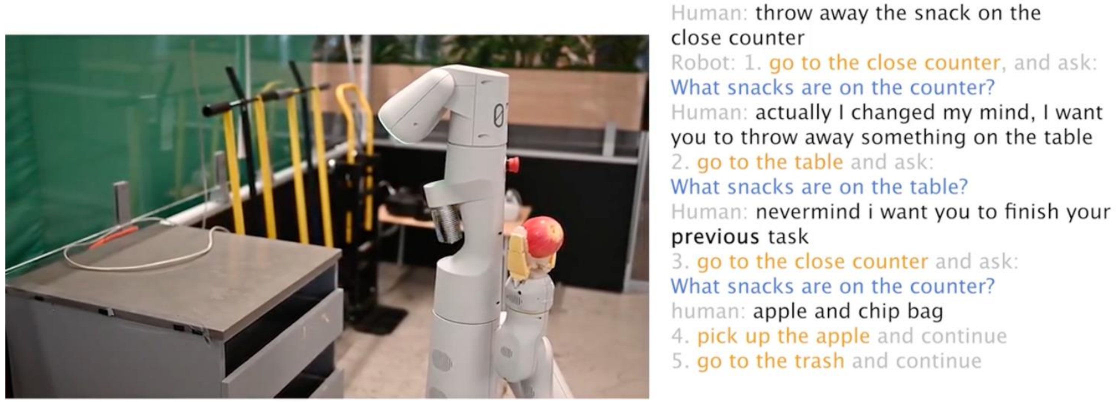

https://innermonologue.github.io/

A Different Path: From pretraining to experience-scaling

➢Effective Experience: Not All Experience is Useful

- **Effective Experience:**
  - Measures how much closed-loop interaction contributes to future reusable skills, adaptation, and regret reduction, rather than how many transitions the agent has stepped.
- **The number of effective experience over a fixed time horizon T is defined as:**

$$N_{\text{eff}}(T) = \sum_{t=1}^{T} w(e_t)$$

counting the interactions that the robots converts into reusable structures for open-ended learning.

this defines the x-axis of the experience scaling law

A Different Path: From pretraining to experience-scaling

➢ **Experience Scaling: the held-out test for open-ended EAI**

- In NLP tasks, scaling laws require a held-out test distribution.
- In embodied AI, we can define the held-out tests as controlled evaluation process:
  - the robot interacts with a hidden, time-varying latent world context c_t
  - while evaluator controls its novelty profile $\boldsymbol{\xi}$.

Not seen by the robot

novelty profile $\boldsymbol{\xi}$: specifies how the OOD shifts change over time.

world context c_t : objects, env dynamics, goal, human preference, sensors, robot embodiments

$$c_{t+1} \sim \mathcal{K}_{\xi}(\cdot \mid c_t)$$

generate continuous task streams to evaluate the EAI models

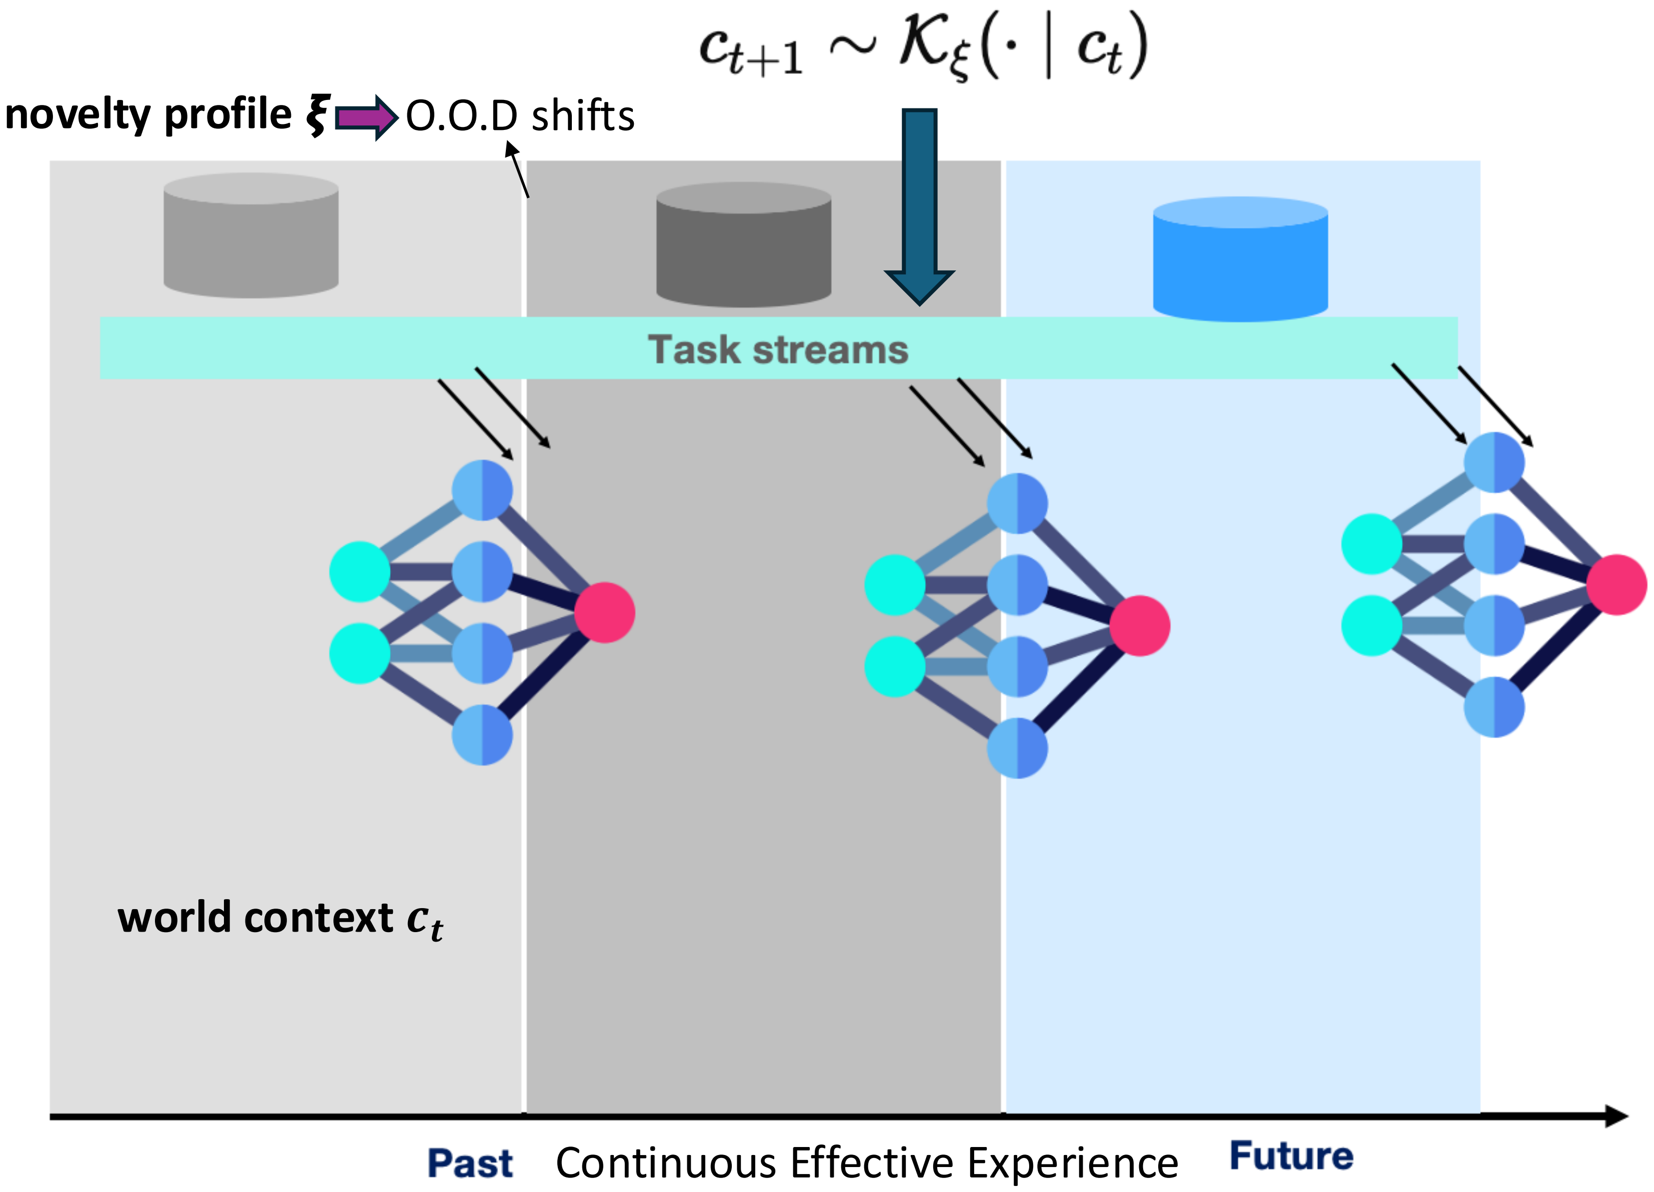

Not seen by the robot

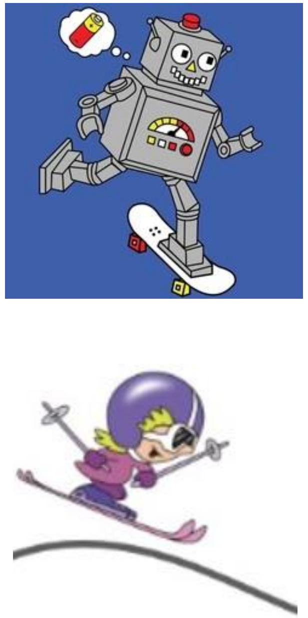

A Different Path: From pretraining to experience-scaling

➢ Experience Scaling: the held-out test for open-ended EAI

- In NLP tasks, scaling laws require a held-out test distribution.
- In embodied AI, we can define the held-out tests as controlled evaluation process:
  - the robot interacts with a hidden, time-varying latent world context c_t
  - while evaluator controls its novelty profile $\boldsymbol{\xi}$.

$$c_{t+1} \sim \mathcal{K}_{\xi}(\cdot \mid c_t)$$

generate continuous task streams to evaluate the EAI models

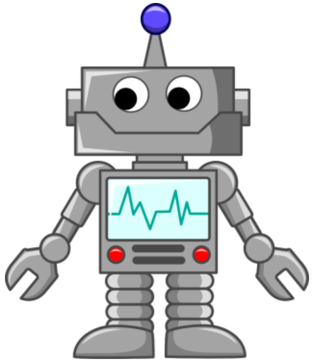

The robot does not observe: $\boldsymbol{c_t}$, $\boldsymbol{\xi}$, task ID

The robot only observes: $h_t = (o_1, a_1, r_1, \ldots, o_t)$

this defines the tests of the experience scaling law

A Different Path: From pretraining to experience-scaling

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- **Deviation-Regret: a task-boundary-free error**
  - Elelimy et al. 2025 defines the deviation regret for a continual learning agent:

$$\underbrace{\rho_T(\phi, \lambda, e)}_{\text{deviation regret}} = \frac{1}{T}\sum_{t=1}^{T}\left(\underbrace{\mathbb{E}\left[\sum_{i=t}^{t+H-1}\gamma^{(i-t)}R_i \middle| \phi(\sigma), H_{t-1}\right]}_{\text{deviation return}} - \underbrace{\mathbb{E}\left[\sum_{i=t}^{t+H-1}\gamma^{(i-t)}R_i \middle| \sigma, H_{t-1}\right]}_{\text{agent return}}\right)$$

  - We follow the same definition, but the tricky part is the $\boldsymbol{\phi}$ (derivations) as it's often unknown.
  - Here, we can treat $\boldsymbol{\phi}$ (derivations)'s outcome as all available policies present to the evaluator.
    - E.g., pretrained VLAs best for this task, world models best for this task, a predefined recovery policy, a predefined controller, etc., which are easy to obtain on the evaluator side.

A Different Path: From pretraining to experience-scaling

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- **Deviation-Regret: a task-boundary-free error**
  - At time t, the agent follows its current policy: $\pi_t = A(h_t)$
  - We compare it against a familiy of availbla edeviation controllers $d \in D_{eval}$,
    - where each d transforms the current policy into an alternative policy: $\pi_t^d = d(\pi_t)$
  - Local H-step return: $J_H(\pi \mid h_t) = \mathbb{E}\left[\sum_{k=0}^{H-1} \gamma^k r_{t+k+1} \mid h_t, \pi\right]$
  - **The deviation-regret over a task stream of length T:**

$$R_H^{\mathcal{D}}(T) = \max_{d \in \mathcal{D}_{\text{eval}}} \sum_{t=1}^{T} \left[J_H(d(\pi_t) \mid h_t) - J_H(\pi_t \mid h_t)\right]$$

return under the best available deviation controller

return under the current policy

A Different Path: From pretraining to experience-scaling

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- **Deviation-Regret: a task-boundary-free error**
  - **Recall the deviation-regret over a task stream of length T:**

$$R_H^{\mathcal{D}}(T) = \max_{d \in \mathcal{D}_{\text{eval}}} \sum_{t=1}^{T} \left[J_H(d(\pi_t) \mid h_t) - J_H(\pi_t \mid h_t) \right]$$

  - **We define expected deviation-regret rate under a fixed evaluation horizon T:**

$$\bar{\rho}_H^{\mathcal{D}_{\text{eval}}}(N_{\text{eff}}) = \mathbb{E}_{\tau \sim \mathcal{D}_{\text{eval}}} \left[\frac{1}{T} R_H^{\mathcal{D}}(\tau; T) \right]$$

Average regret by NOT using the best systematic intervention in $\mathcal{D}_{\text{eval}}$

Did the robot continually learn well in the past T time window?

this defines the y-axis of the experience scaling law

➢**Experience Scaling:** how to measure if the robot learns well in open-ended worlds?

Expected Deviation-Regret Rate on held-out OOD Tasks $\bar{\rho}_H^{\mathcal{D}_{\text{eval}}}(N_{\text{eff}}) = \mathbb{E}_{\tau \sim \mathcal{D}_{\text{eval}}}\left[\frac{1}{T} R_H^{\mathcal{D}}(\tau; T)\right]$

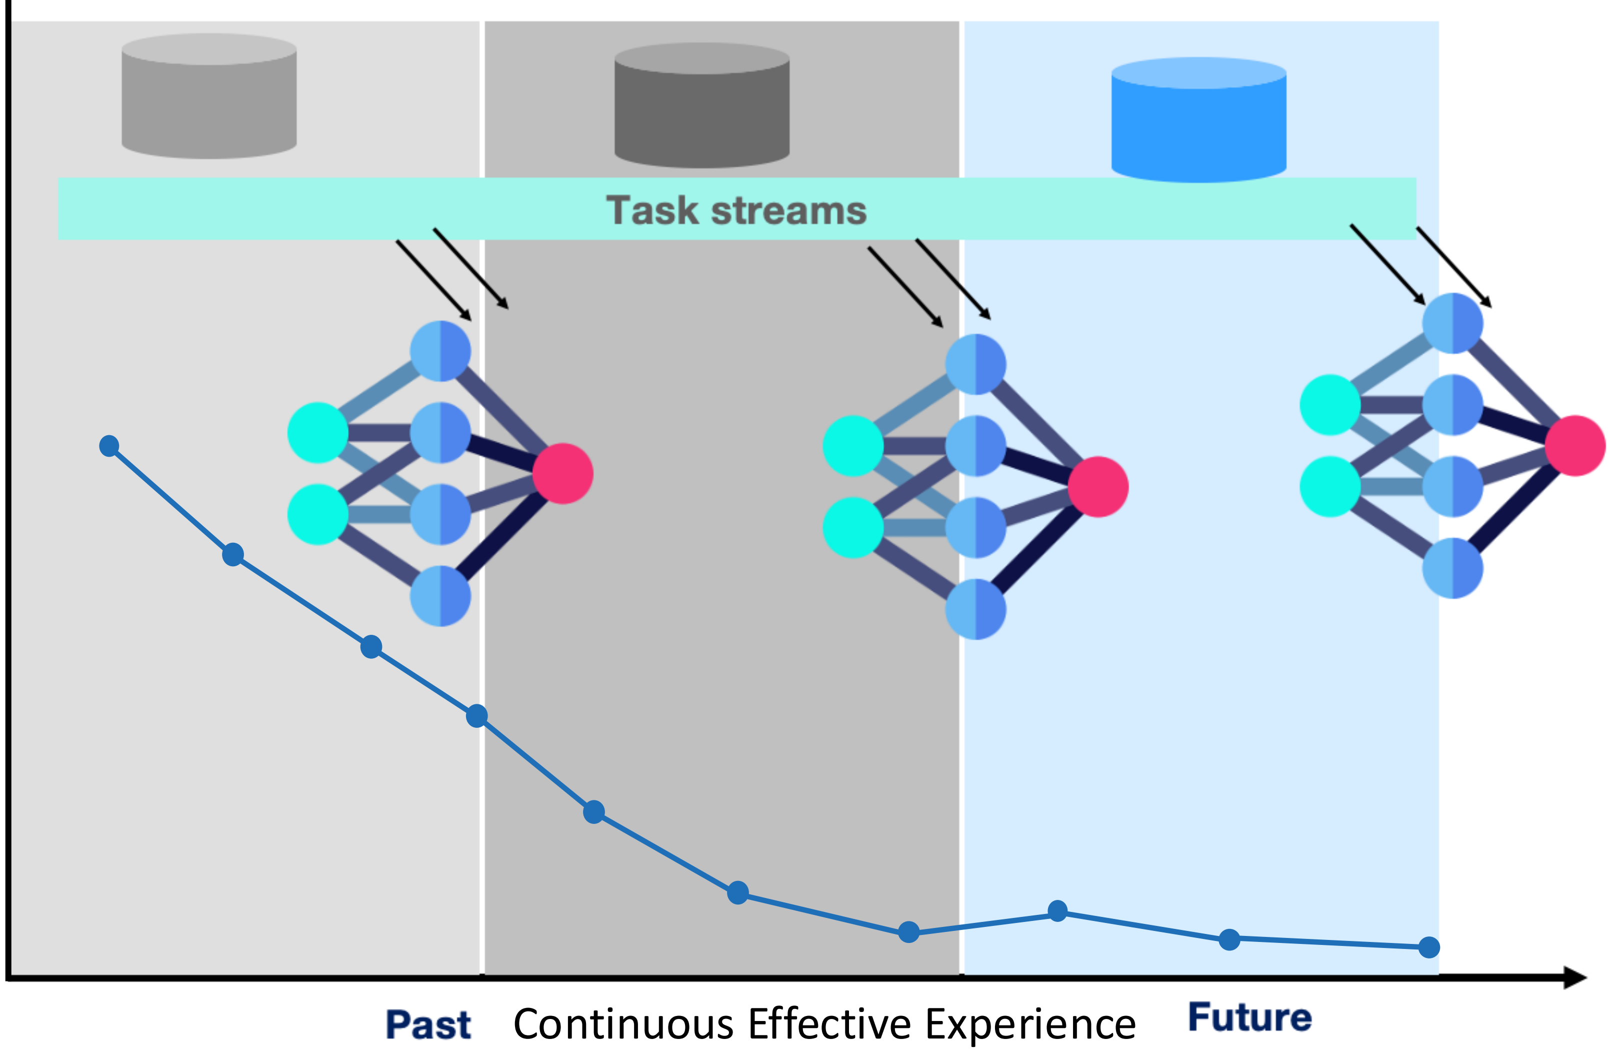

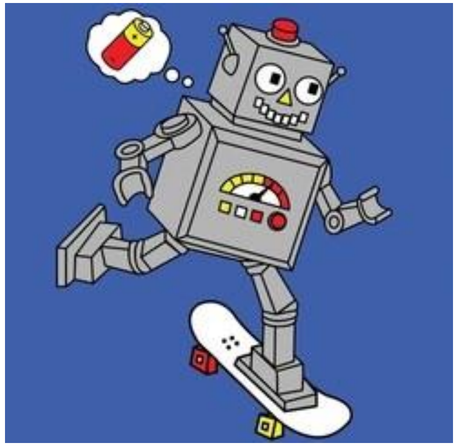

$$N_{\text{eff}}(T) = \sum_{t=1}^{T} w(e_t)$$

Cumulative Effective Experience N_eff

➢**Experience Scaling: how to measure if the robot learns well in open-ended worlds?**

- Suppose a strong Experience-scaling embodied agent (ESEA) that **solves the open-ended embodied intelligence problem**, we aim to see the experience scaling as shown below.

➢ Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- Suppose a strong Experience-scaling embodied agent (ESEA) that **solves the open-ended embodied intelligence problem**, we aim to see the experience scaling as shown below.

Expected Deviation-Regret Rate on held-out OOD Tasks

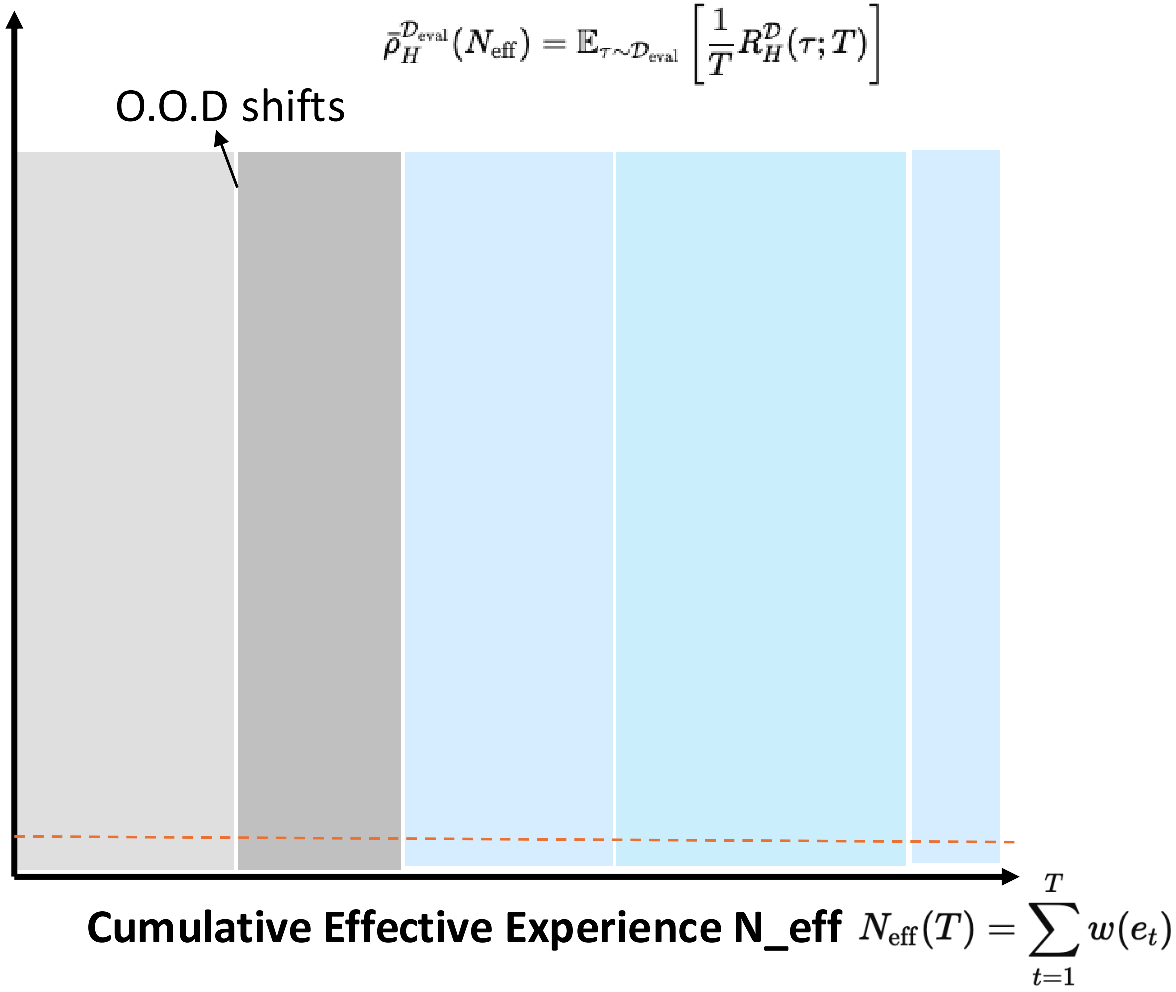

Cumulative Effective Experience

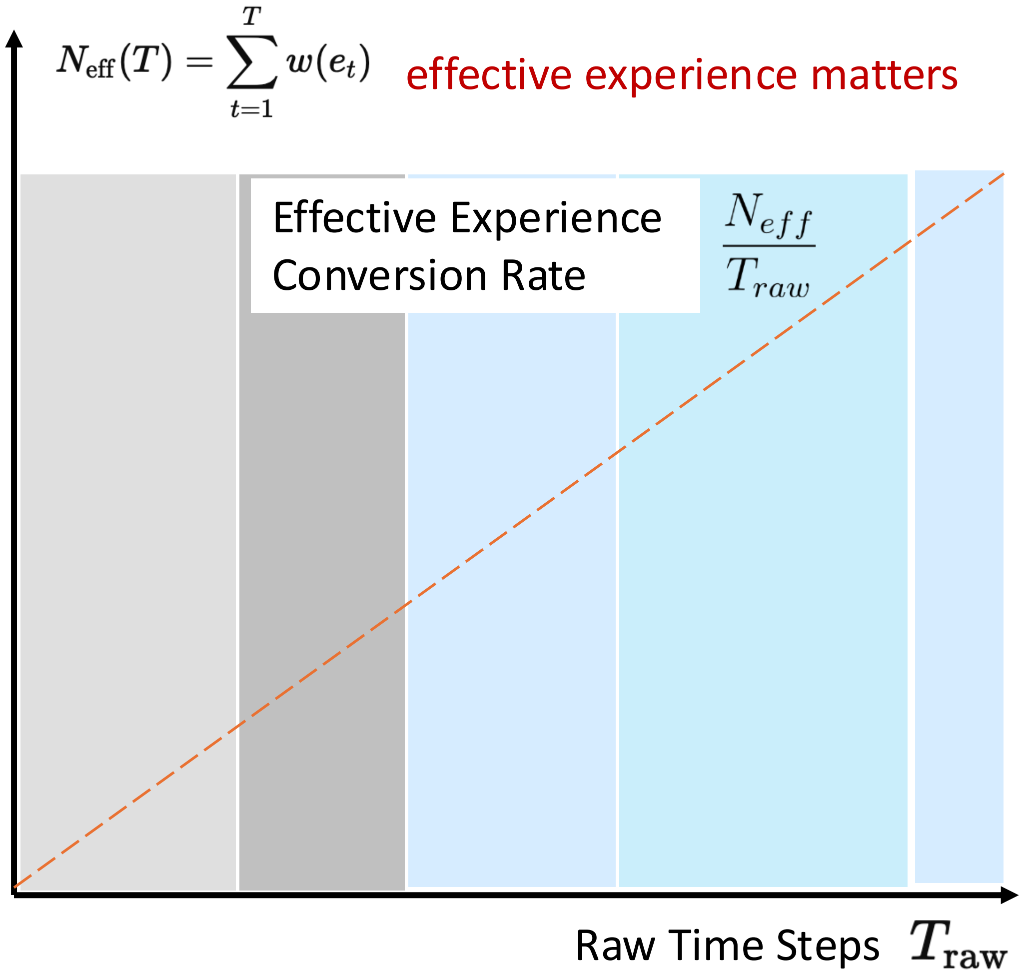

**Conceptual illustration only —since no existing benchmarks are at present.

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- Suppose a strong Experience-scaling embodied agent (ESEA) that **solves the open-ended embodied intelligence problem**, we aim to see the experience scaling as shown below.

Expected Deviation-Regret Rate on held-out OOD Tasks

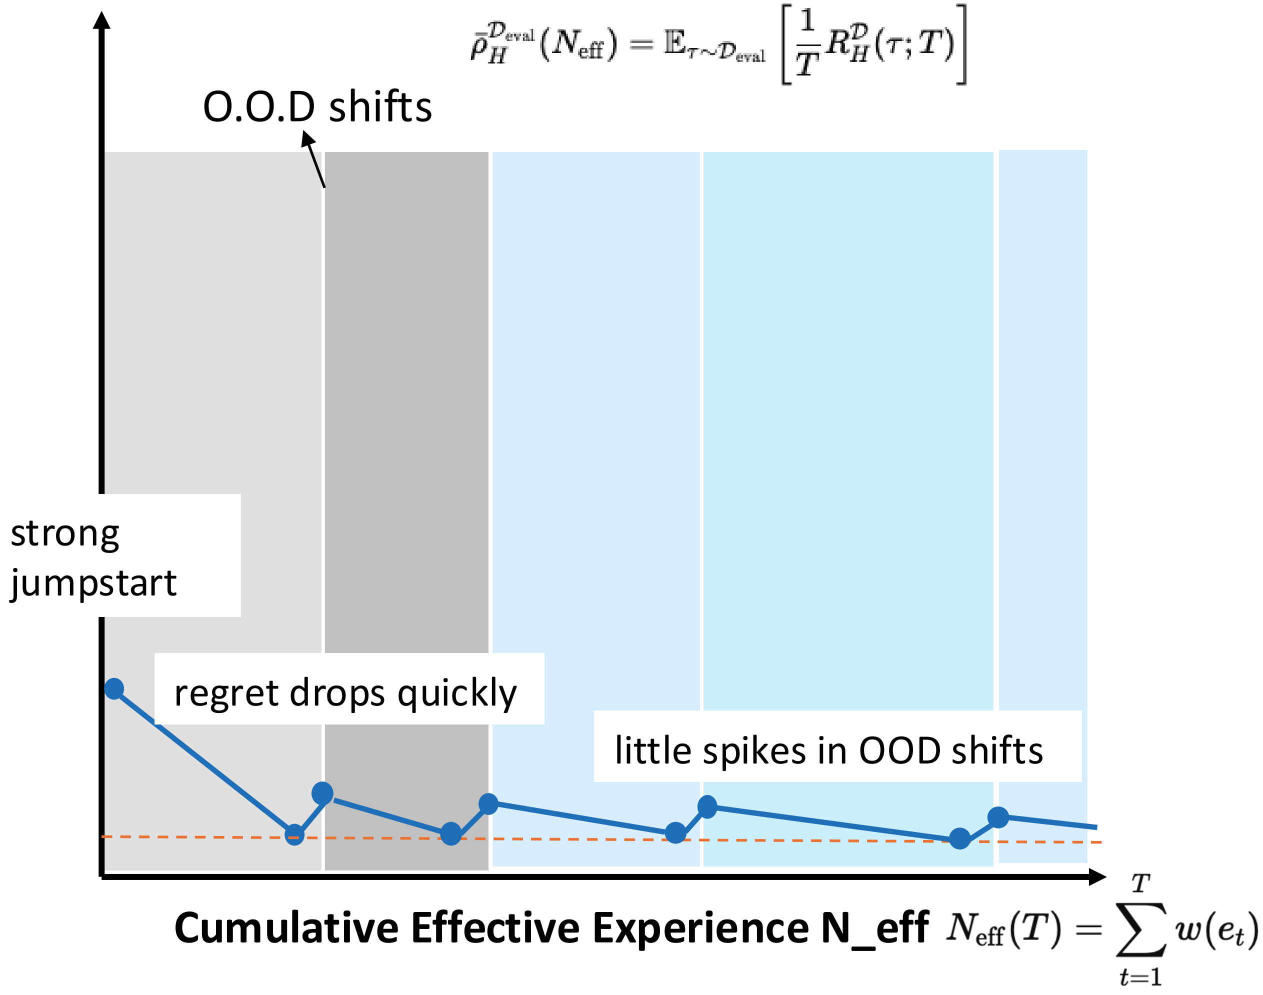

Cumulative Effective Experience

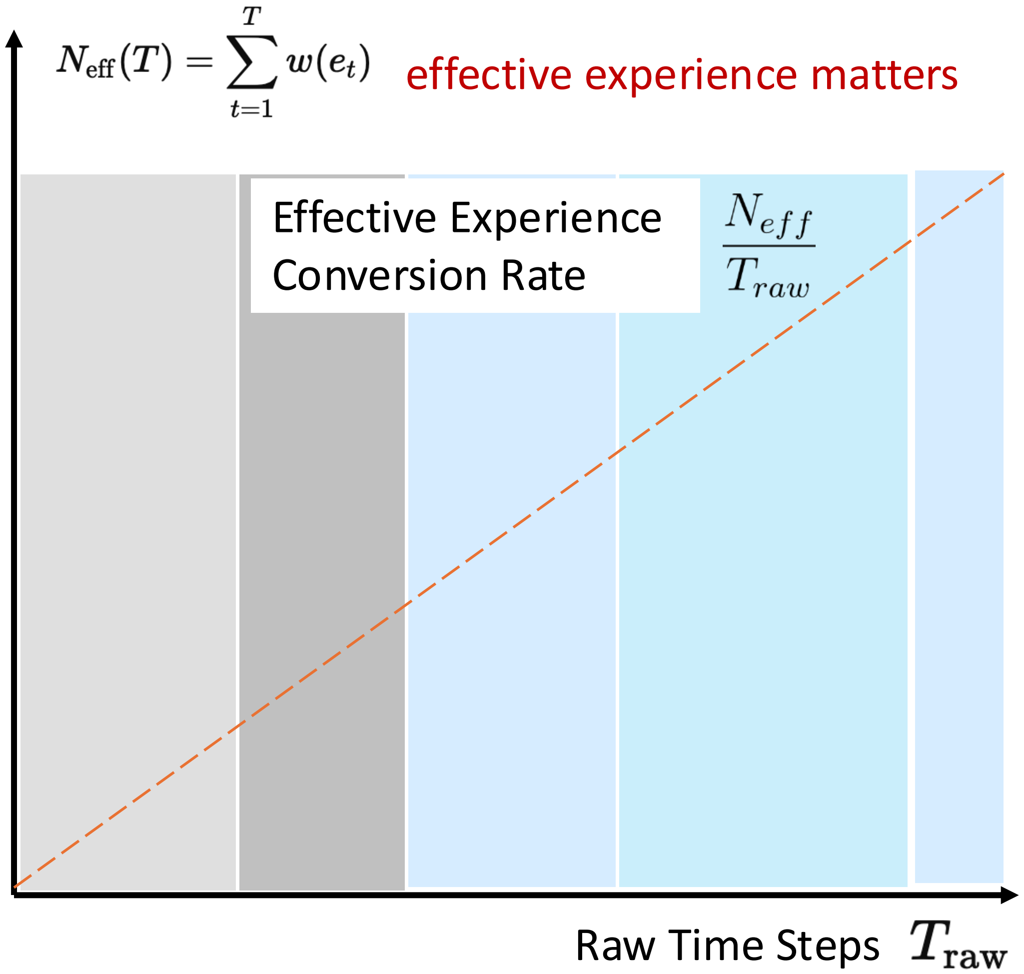

**Conceptual illustration only —since no existing benchmarks are at present.

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- Suppose a strong Experience-scaling embodied agent (ESEA) that **solves the open-ended embodied intelligence problem**, we aim to see the experience scaling as shown below.

Expected Deviation-Regret Rate on held-out OOD Tasks

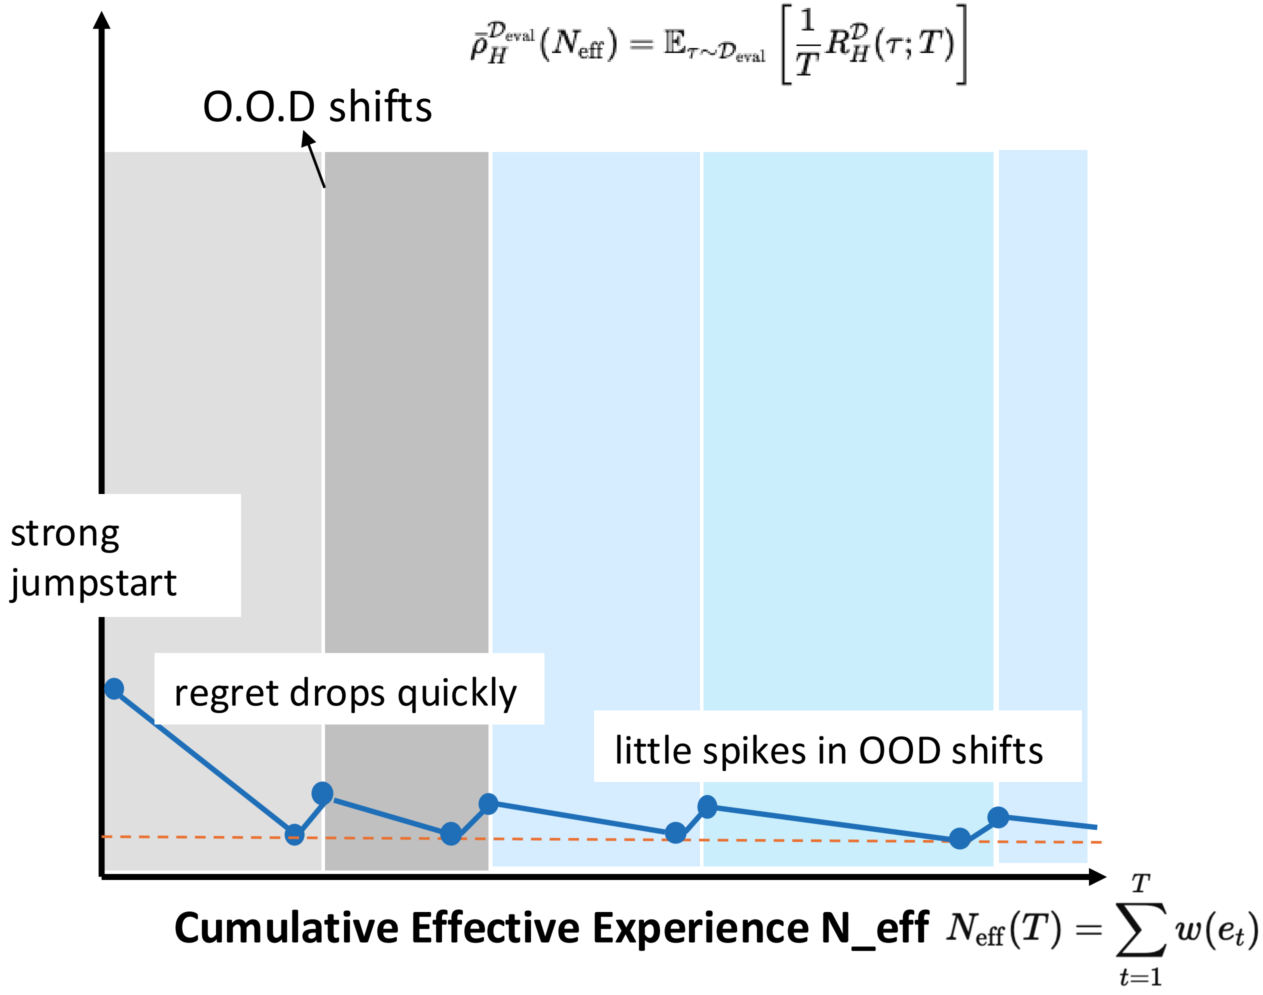

Cumulative Effective Experience

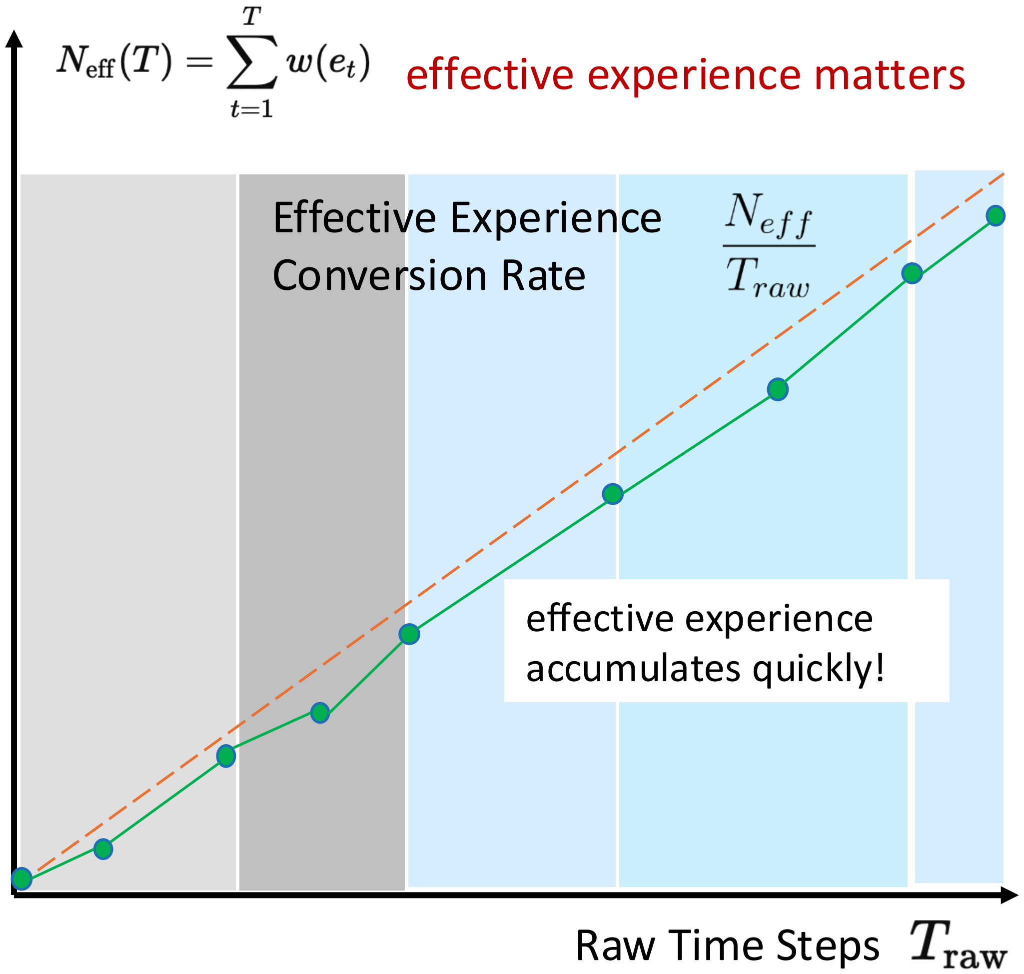

**Conceptual illustration only —since no existing benchmarks are at present.

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- How about a static pretrained EAI model?

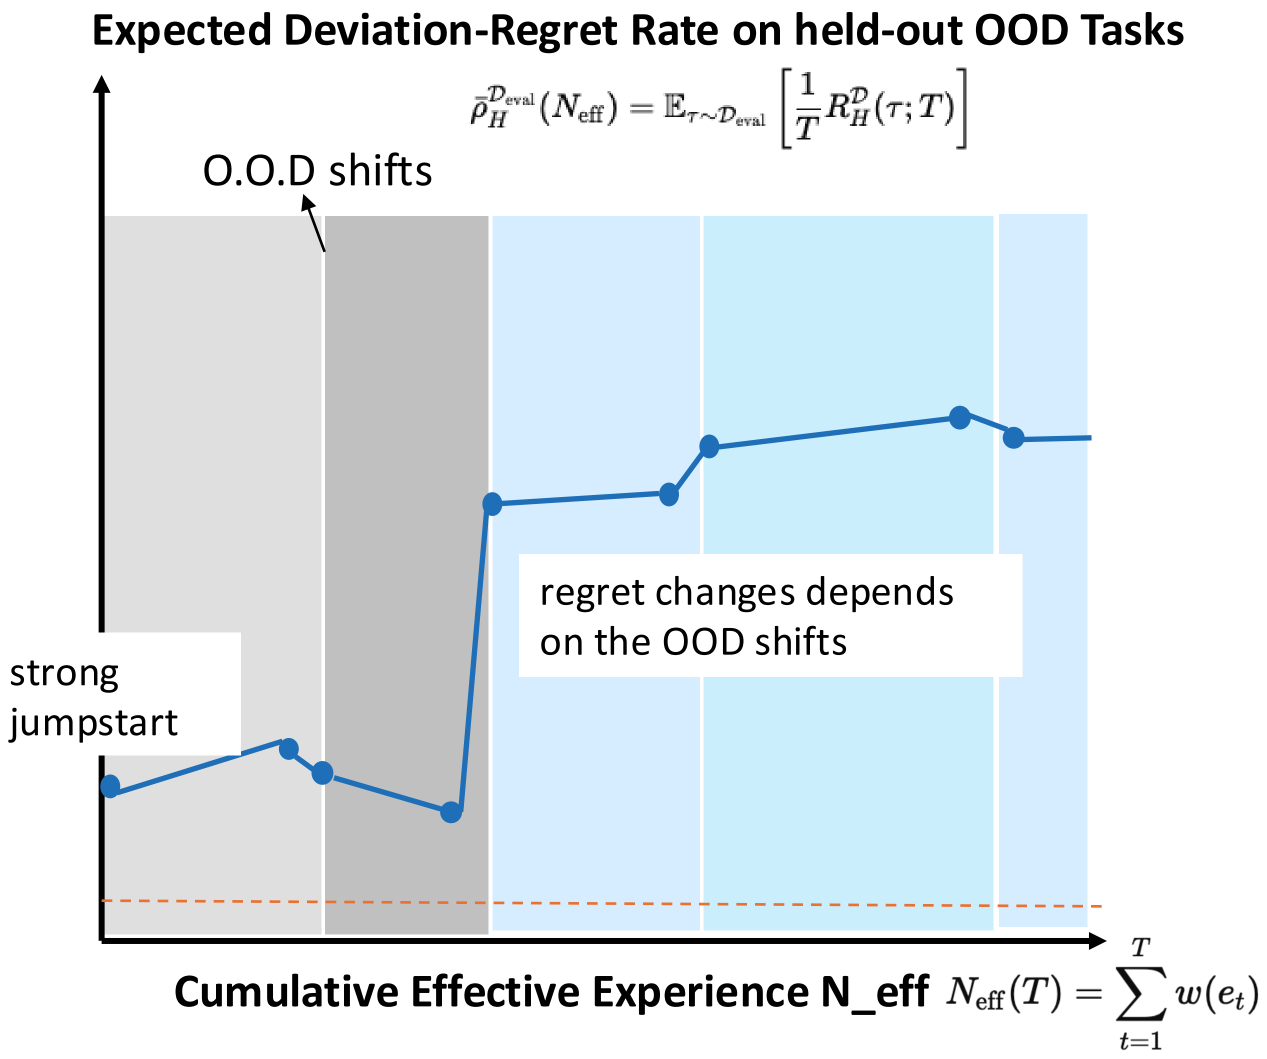

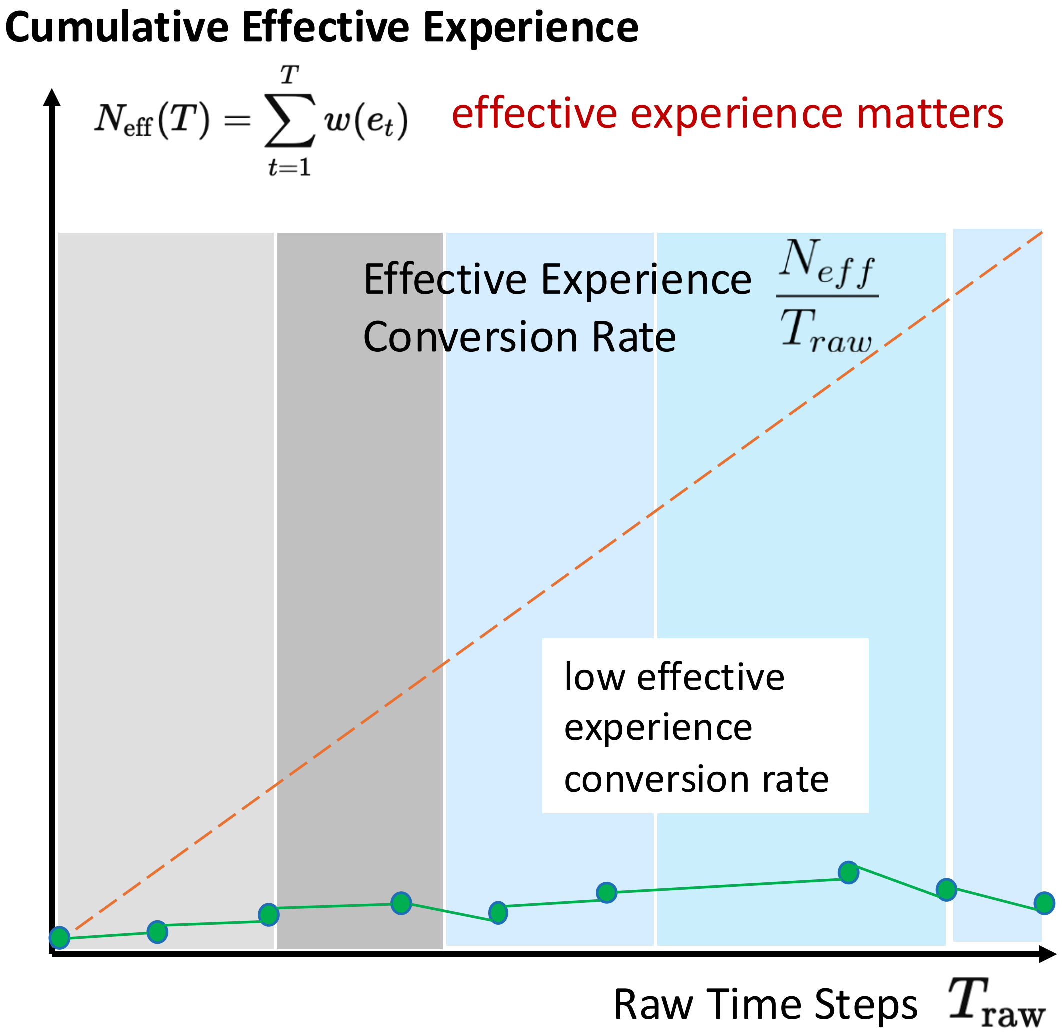

**Conceptual illustration only —since no existing benchmarks are at present.

➢Experience Scaling: how to measure if the robot learns well in open-ended worlds?

- How about a static pretrained EAI model?

Expected Deviation-Regret Rate on held-out OOD Tasks

Cumulative Effective Experience

O.O.D shifts

effective experience matters

Effective Experience Conversion Rate

$$\frac{N_{eff}}{T_{raw}}$$

- **How about EAI post training using RL?**
- **Do we have such benchmarks now?**

ARC Prize Foundation, 2026

close, but not sufficient…

strong jumpstart

low effective experience conversion rate

Cumulative Effective Experience N_eff $N_{\text{eff}}(T) = \sum_{t=1}^{T} w(e_t)$

Raw Time Steps T_{raw}

**Conceptual illustration only —since no existing benchmarks are at present.

A Different Path: From pretraining to experience-scaling

➢If static pretraining is NOT sufficient, what are alternative pathways?

- **Will post-training (SFT & RL) be the rescue?**
  - SFT and RL post-training are not a satisfactory general solution either, because they are difficult to scale across continually changing OOD shifts.
- Can we think about the problem in this way?
  - Pretraining gives strong priors, we agree on that.
  - But what is missing is a way for the robot's embodied capability to keep growing through interaction with the unbounded world under bounded resources (compute, energy consumption).
  - The missing axis here is experience!

image source: dreamstime.com

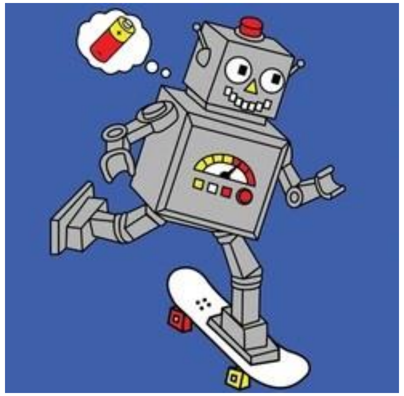

A Different Path: From pretraining to experience-scaling

➢Artifacts, external memory, and the role of pretrained foundation models

- Martin et al. 2026 (Mar, Amii Talk): environmental artifacts can reduce the amount of internal information needed to represent history and act competently.
- **Following this idea, pretrained foundation models can be interpreted as external cognitive artifacts.**

agent + artifacts + environments

- VLM: semantic and perceptual library for the robot
- VLA: an action prior and policy initializer.
- World model: predictive artifact for sub-task proposal testing.

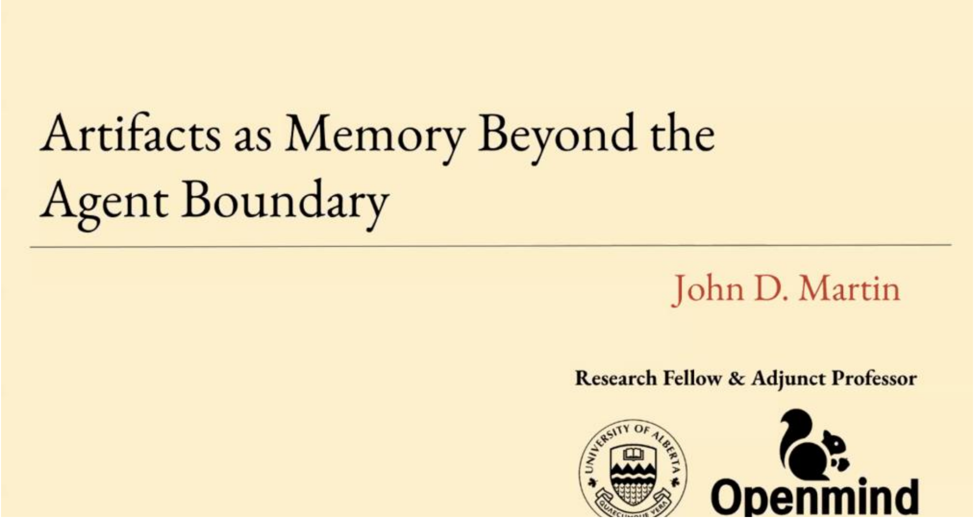

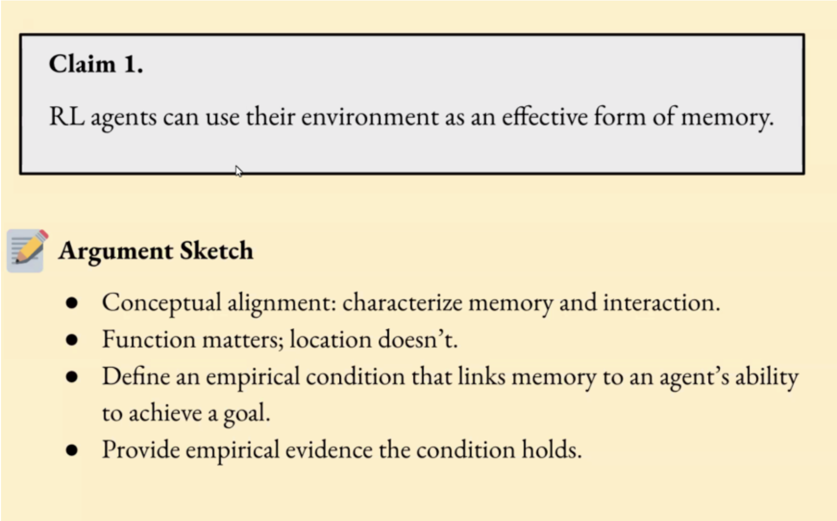

A Different Path: From pretraining to experience-scaling

➢My sketch draft: the Experience Scaling Embodied Agent

- Experience-scaling embodied agent (ESEA) is an agent architecture that converts embodied interaction into reusable skills under resource constraints.

onboard policy and representation

memory and skill library

$$\mathcal{A} = (\theta, \mathcal{M}, \mathcal{E}, \mathcal{R})$$

external artifacts (pretrained LLMs, VLMs, VLAs)

resource bounds (computation, energy)

A Different Path: From pretraining to experience-scaling

➢ **Open Research Questions:** **we can redefine EAI research!**

OVERALL RESEARCH QUESTION

RQ0. How can we build a unified embodied foundation architecture that integrates pretrained foundation models as memory-artifact components and continuously turn embodied experience into reusable skills, adaptive memory, and reduced deviation regret?

A Different Path: From pretraining to experience-scaling

➢Open Research Questions: we can redefine EAI research!

SPECIFIC RESEARCH QUESTION

- ❑ Q1: How can pretrained foundation models, policies, skills, and memory structure be integrated into a unified embodied foundation architecture?
- ❑ Q2: How should foundation models be formalized as memory-artifact components inside embodied agents as compressed memory structures?
- ❑ Q3: How can embodied experience update the agent's **memory** structure, skill library, and learned representations?
- ❑ Q4: How can agents **discover reusable skills**, options, subgoals, and abstractions under memory-artifact components and bounded resource limits?
- ❑ Q5: How can experience scaling laws be validated in simulated and real-world benchmarks?

The Pursuit of General-Purpose ~~Robots~~ 👉 Tools

I prefer building general-purpose tools for humans.

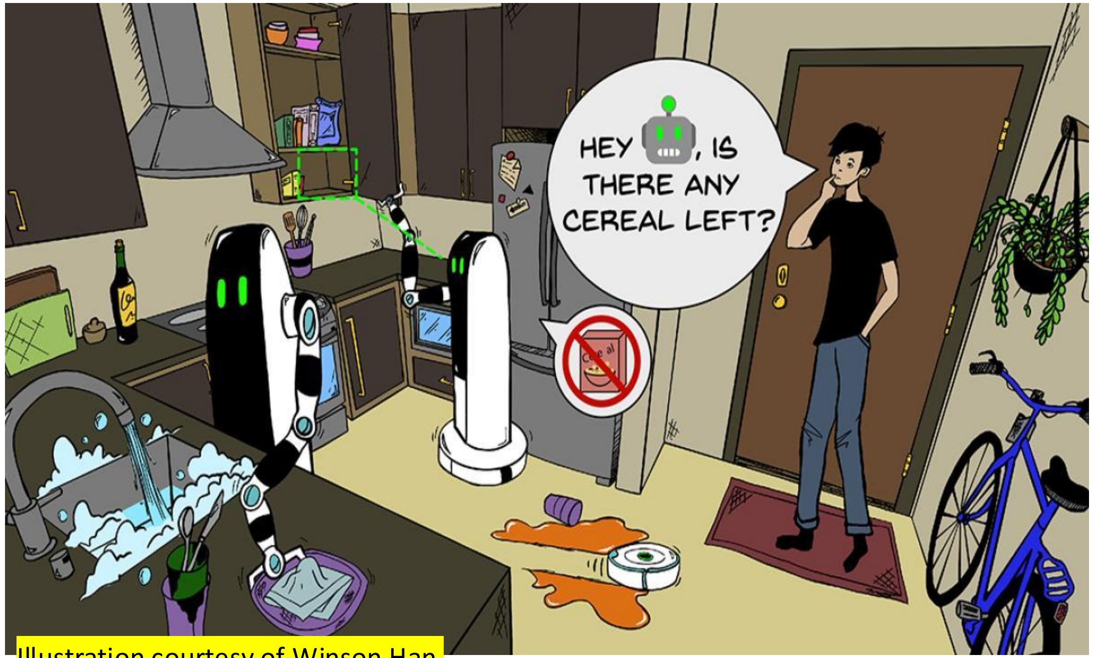

Illustration courtesy of Winson Han.

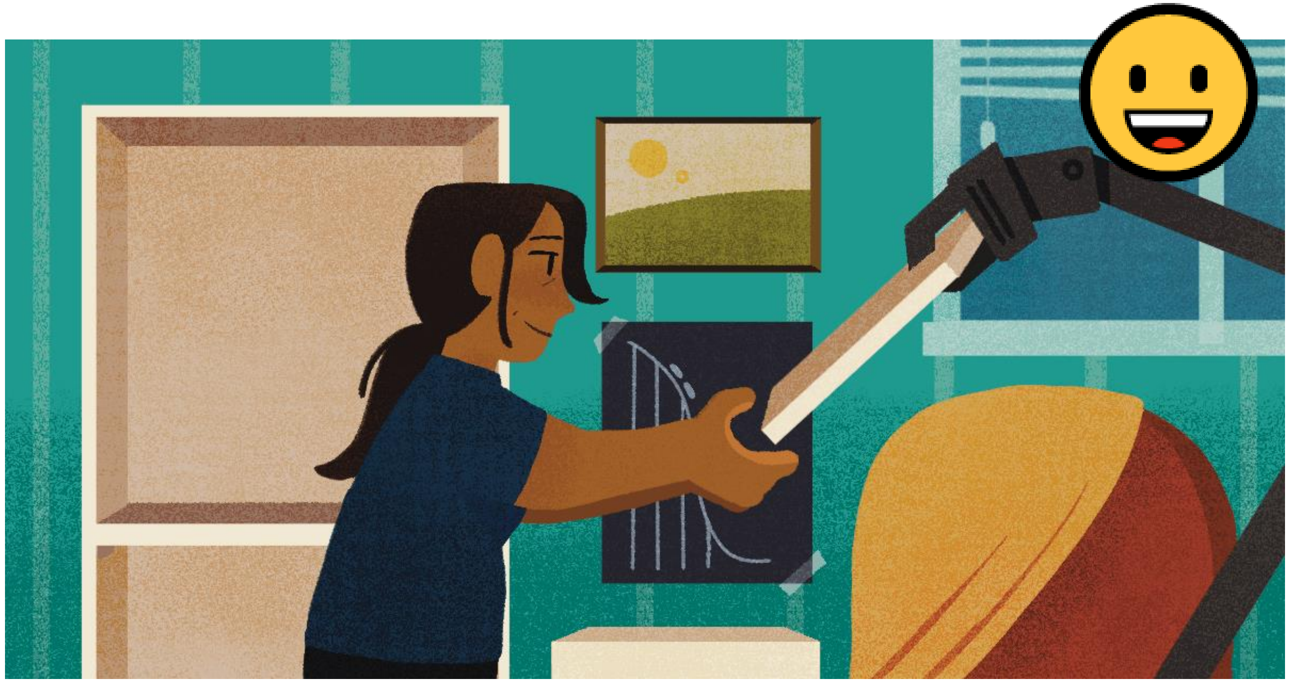

https://engineering.oregonstate.edu/all-stories/beyond-chatgpt-embodied-language-understanding

https://viterbischool.usc.edu/news/2021/06/need-help-building-ikea-furniture-this-robot-learns-your-preferences-and-lends-a-hand/

Hiring:

- I am recruiting MSc and PhD students.
- If you are interested in this research direction, welcome to apply to my lab!

Human-Centered Autonomy Lab
AI AND ROBOTICS FOR SOCIAL GOOD

amii

UNIVERSITY OF ALBERTA

ALBERTA INNOVATES

INNOVATION
Canada Foundation for Innovation
Fondation canadienne pour l'innovation

NSERC CRSNG

CIFAR

Dr. Jun Jin
jjin5@ualberta.ca
www.ece.ualberta.ca/~jjin5/

Thank you!

backup slides

A Different Path: From pretraining to experience-scaling

➢Why current continual RL is still not ready?

- **Current continual RL usually studies a “cleaner” problem than the one real embodied agents face.**
  - Episodic evaluation vs. never-ending interaction.
  - Expected return vs. regret
  - Forgetting as a bug vs. forgetting as a necessity
  - Convergence to optimal artifacts vs. persistent adaptation
  - Scaling model capacity vs. surviving under fixed resources

Roadmap of EAI Models

➢Phase 2: 0.5 End-to-End Models / Vision-Language-Action (VLA) models

(2) Diffusion-based Action Expert ➡ **much better high DOF / continuous control performance**

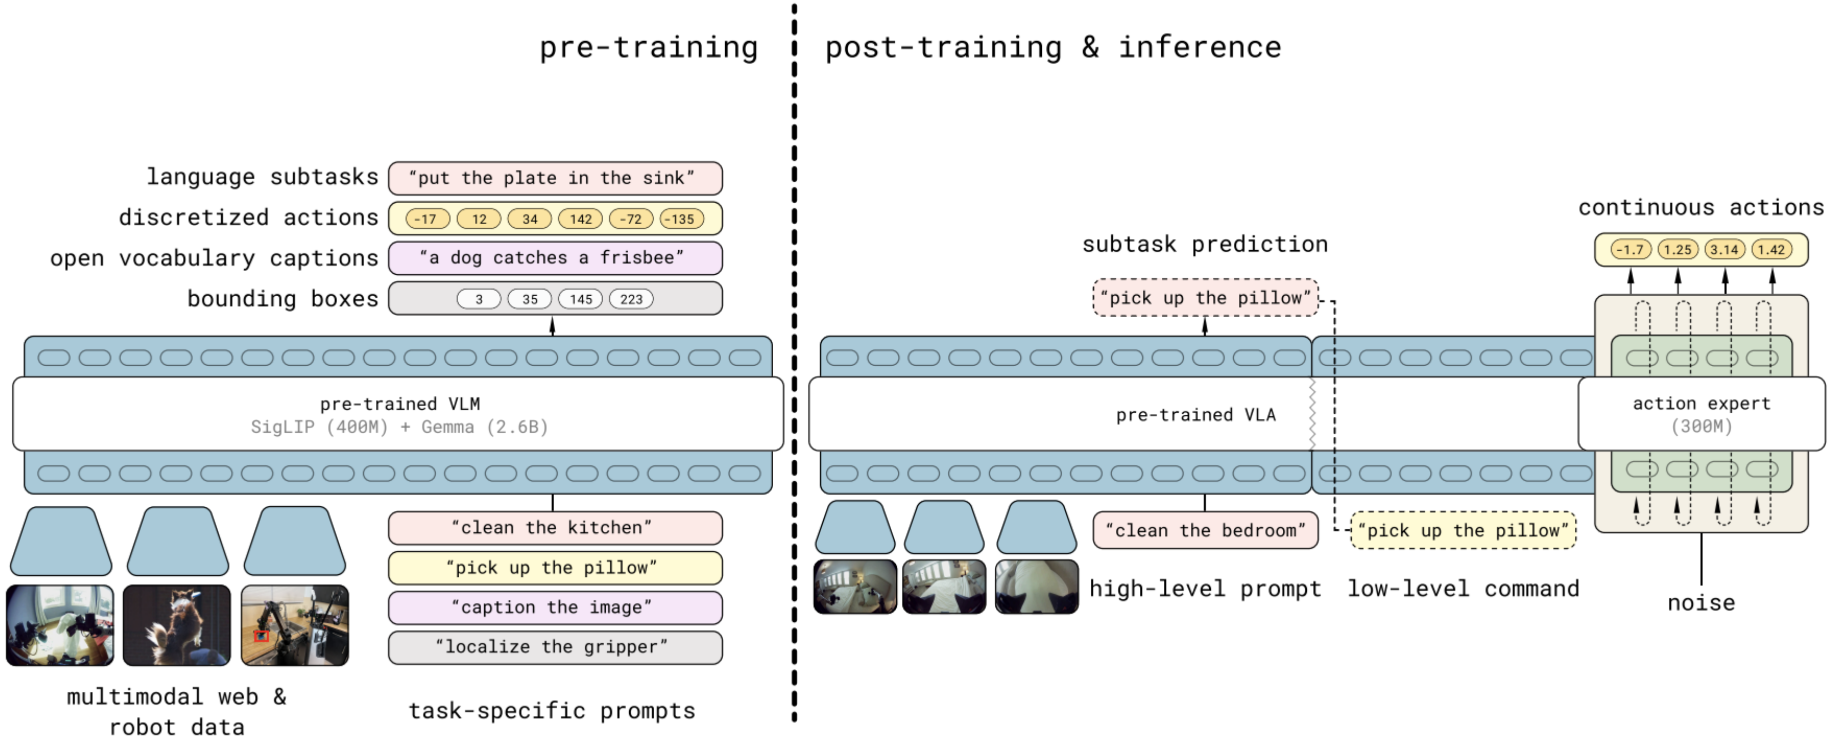

Pi-0.5, Black et al. 2024