Robotic Learning with Large Datasets:
Towards General-Purpose Models for All Robots

Sergey Levine

UC Berkeley

How AI used to work

segmentation model

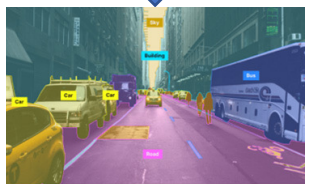

classification model

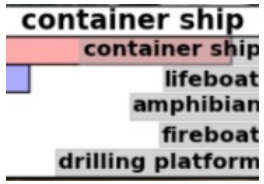

captioning model

A group of people shopping at an outdoor market.

What is the mustache made of?

visual QA model

bananas

"Horrible services. The room was dirty and unpleasant. Not worth the money."

sentiment model

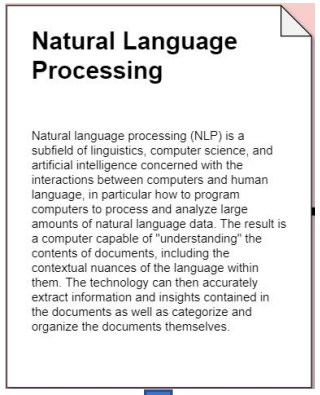

summarization model

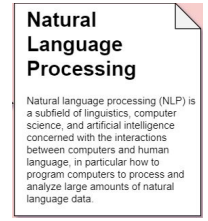

How AI works now

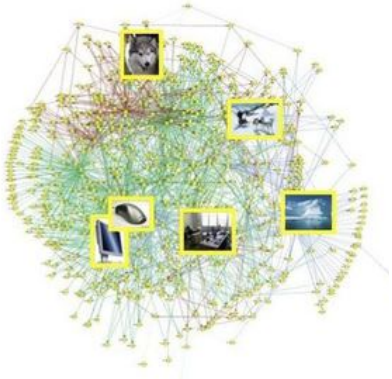

giant un/self-supervised pretrained model

“foundation model”

finetuning

segmentation model

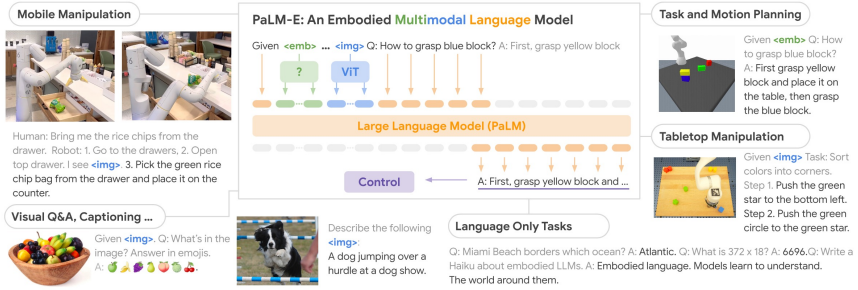

Dreiss et al., PaLM-E: An Embodied Multimodal Language Model. 2023.

How robotic learning works now

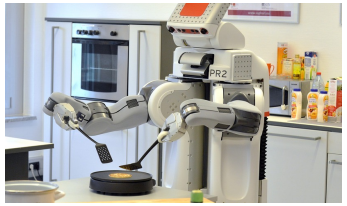

PR2 pancake model

WAM pancake model

UR10 box picking model

lipstick robot model

mousetrap… hand… shake… model??

How robotic learning will work in the future

giant un/self-supervised pretrained **robot** model

"robot foundation model"

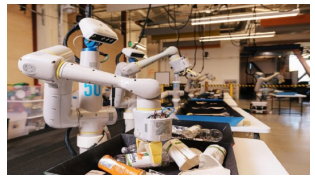

finetuning

trash sorting model

What would it take to build such a general-purpose robot model?

How does this look for robots?

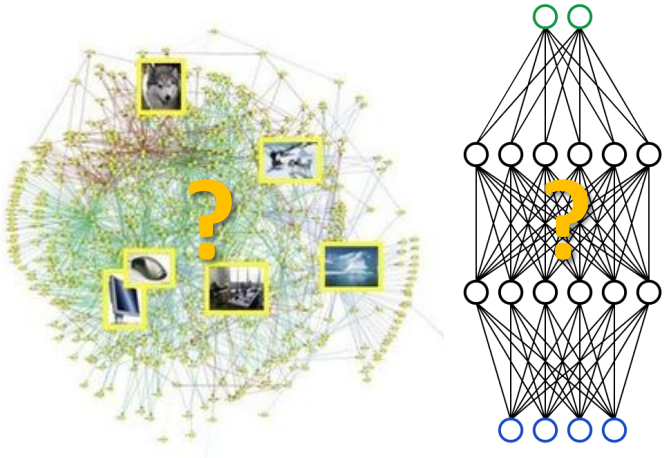

The data needs to…

- Be diverse enough to support my (new) robot
- Cover a range of environments
- Cover a range of behaviors

The model needs to…

- Be able to train on all this data
- Be flexible enough to repurpose to new tasks
- Perform at least some tasks in zero shot?

Bridge data

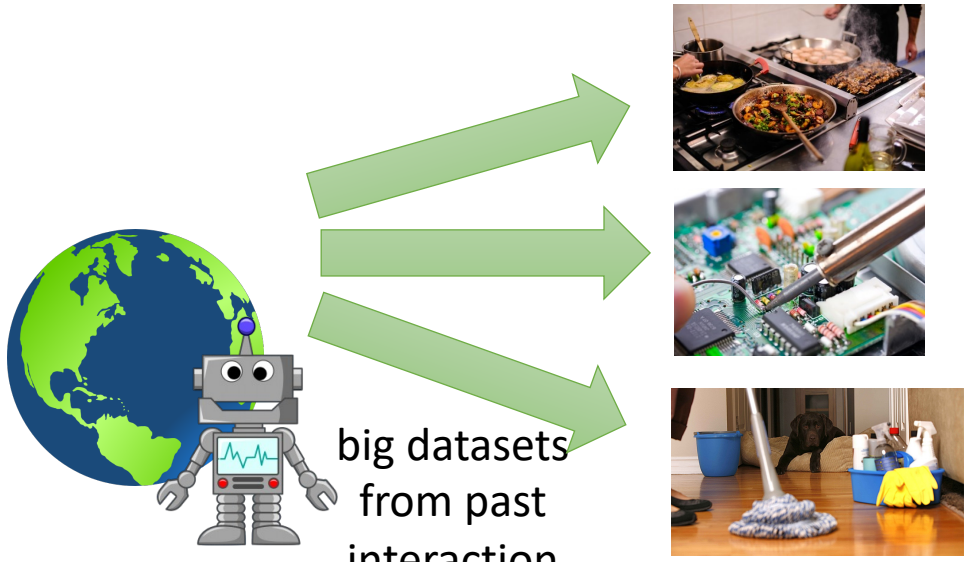

- 33,000+ demonstrations
- 20+ environments
- 100+ tasks
- Designed to be reusable by other researchers in **new domains** and for **new tasks**

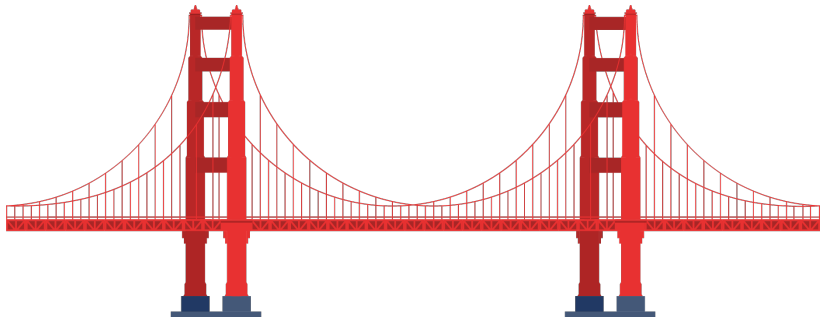

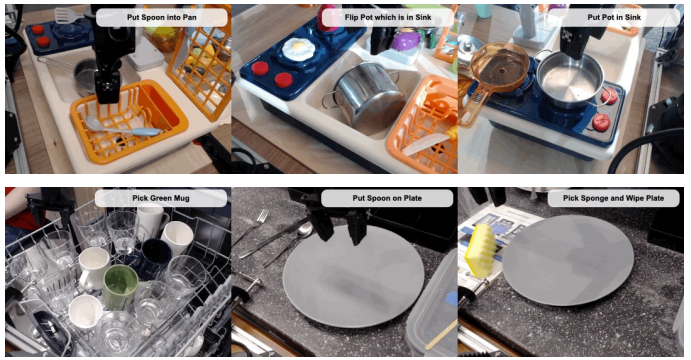

Bridge Data: https://rail-berkeley.github.io/bridgedata/
Led by Homer Walke, Frederik Ebert, with many others

Pretraining on bridge data with offline RL (PTR)

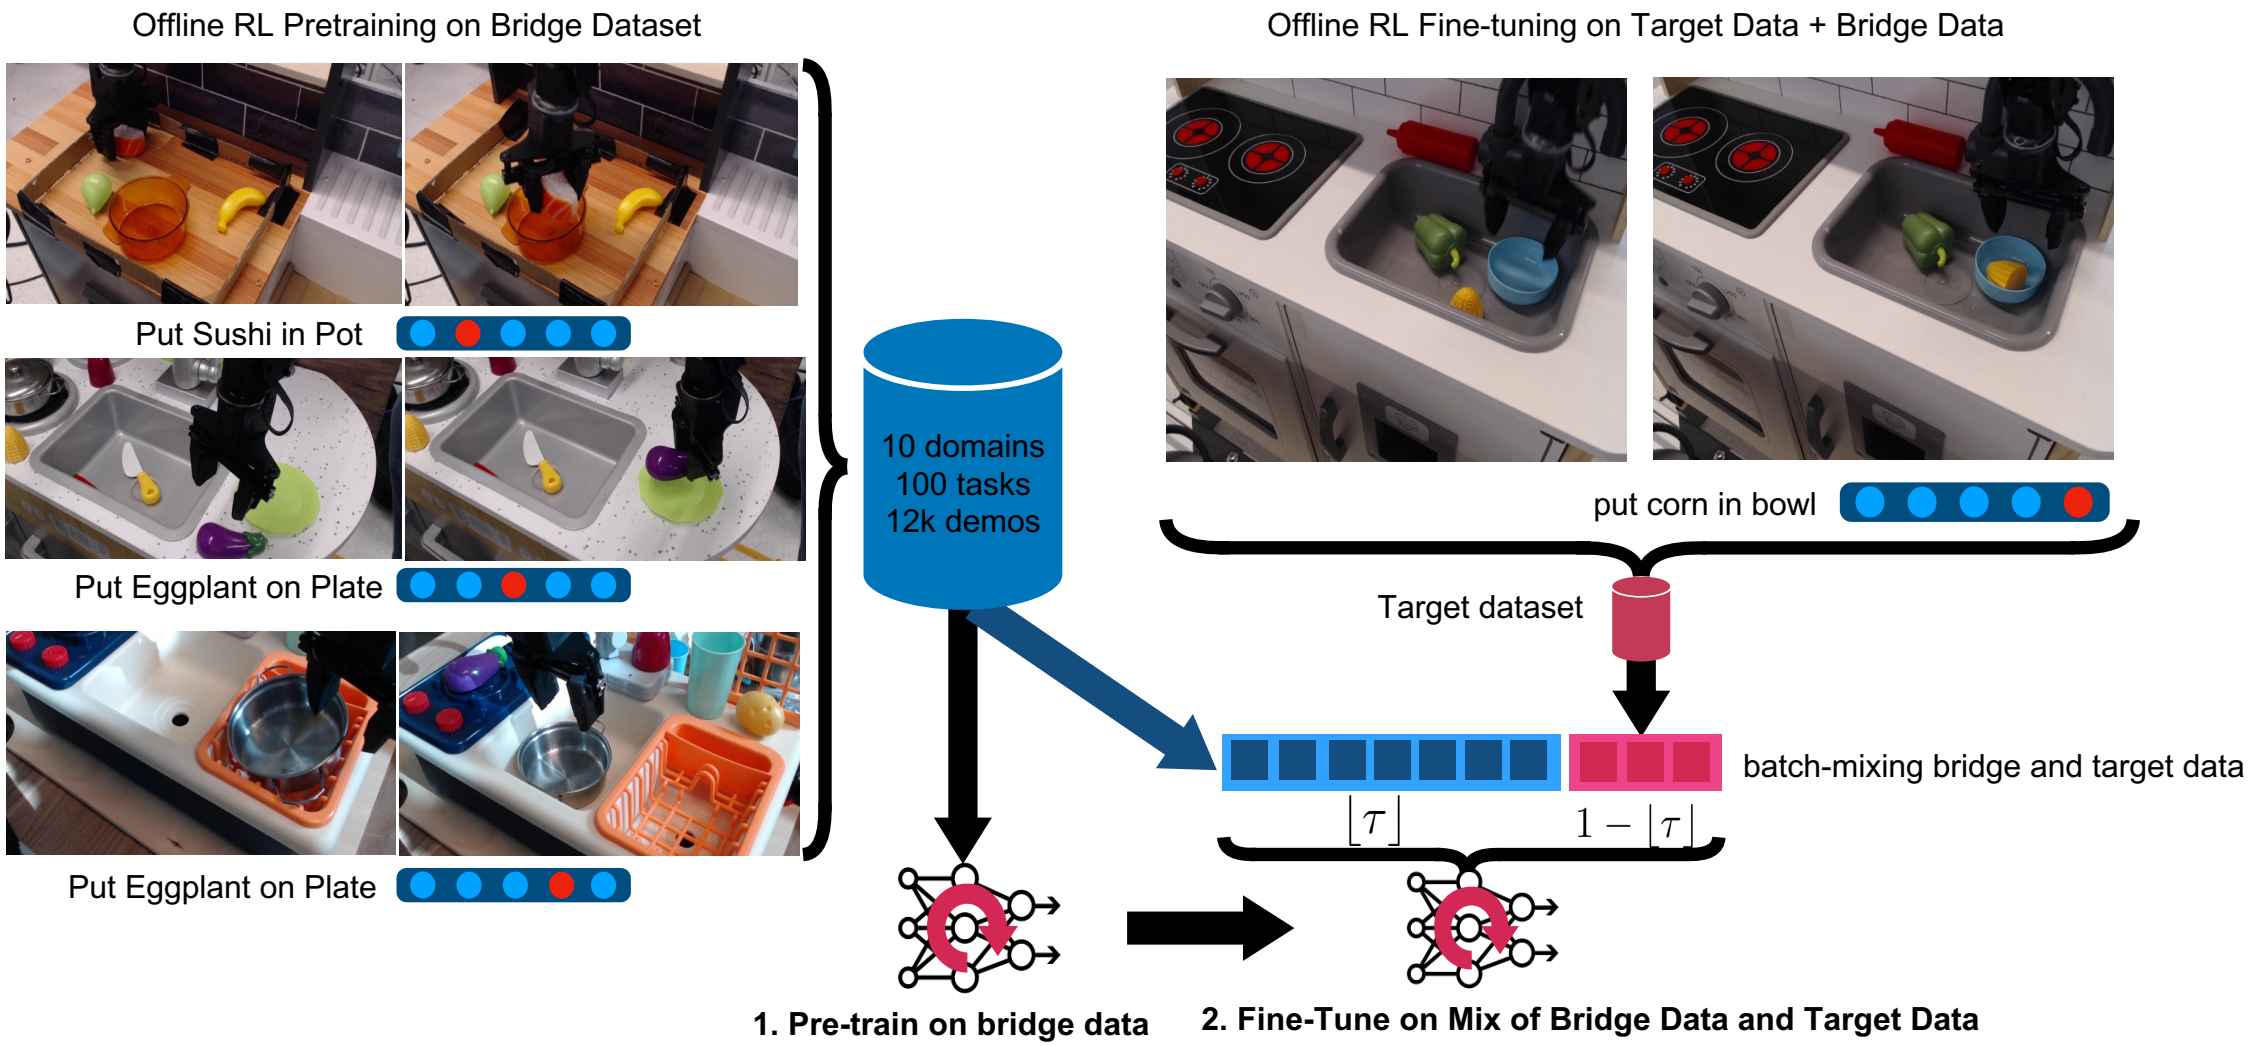

Aviral Kumar, Frederik Ebert, Yanlai Yang, Anikait Singh, Chelsea Finn, Sergey Levine. **Pre-Training for Robots: Offline RL Enables Learning New Tasks from a Handful of Trials.** 2022.

PTR (Pre-Training for Robots) Results

learning **entirely new** skills after pretraining on the bridge dataset

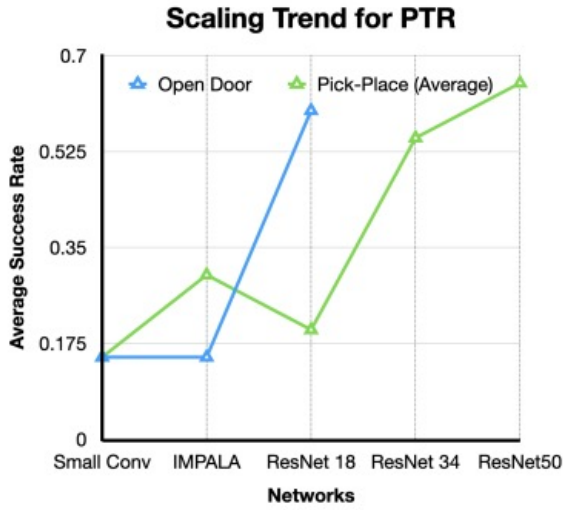

		BC finetuning			Joint training		Target data only		Pre-train. rep. + BC finetune	
Task	**PTR (Ours)**	**BC (fine.)**	**Autoreg. BC**	**BeT**	**COG**	**BC**	**CQL**	**BC**	**R3M**	**MAE**
Take croissant from metal bowl	**7/10**	3/10	5/10	1/10	4/10	4/10	0/10	1/10	1/10	3/10
Put sweet potato on plate	**7/20**	1/20	1/20	0/20	0/20	0/20	0/20	0/20	0/20	1/20
Place knife in pot	**4/10**	2/10	2/10	0/10	1/10	3/10	3/10	0/10	0/10	0/10
Put cucumber in pot	**5/10**	0/10	1/10	0/10	2/10	1/10	0/10	0/10	0/10	0/10

Aviral Kumar, Frederik Ebert, Yanlai Yang, Anikait Singh, Chelsea Finn, Sergey Levine. **Pre-Training for Robots: Offline RL Enables Learning New Tasks from a Handful of Trials.** 2022.

An unsupervised subgoal planning framework

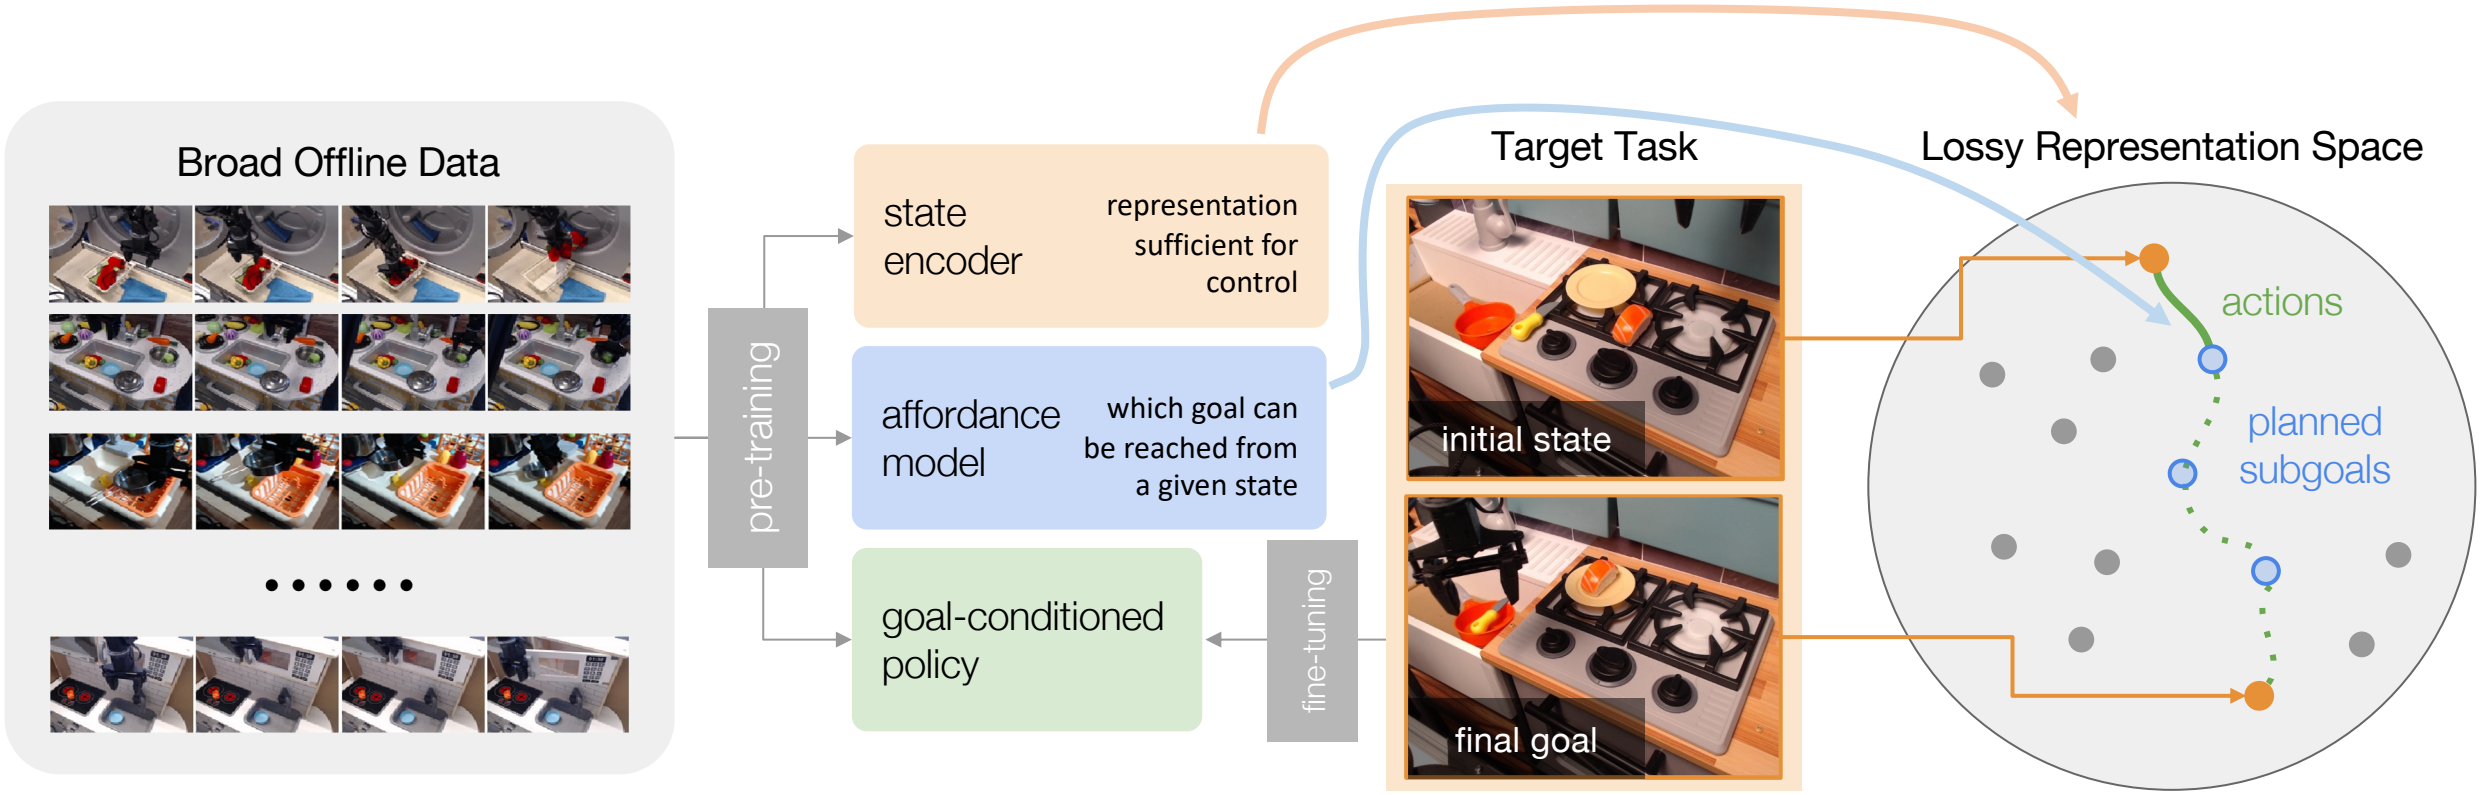

Fang, Yin, Nair, Walke, Yan, Levine. **Generalization with Lossy Affordances: Leveraging Broad Offline Data for Learning Visuomotor Tasks**. 2022.

Task 1: **Move colander onto stove and drop object into colander**

Model-Free
0.0% Success Rate

FLAP (Target Data Only)
25.0% Success Rate

75.0% Success Rate

FLAP (Broad Data)

Fang, Yin, Nair, Walke, Yan, Levine. **Generalization with Lossy Affordances: Leveraging Broad Offline Data for Learning Visuomotor Tasks**. 2022.

Massively **multi-task** policies: RT-1

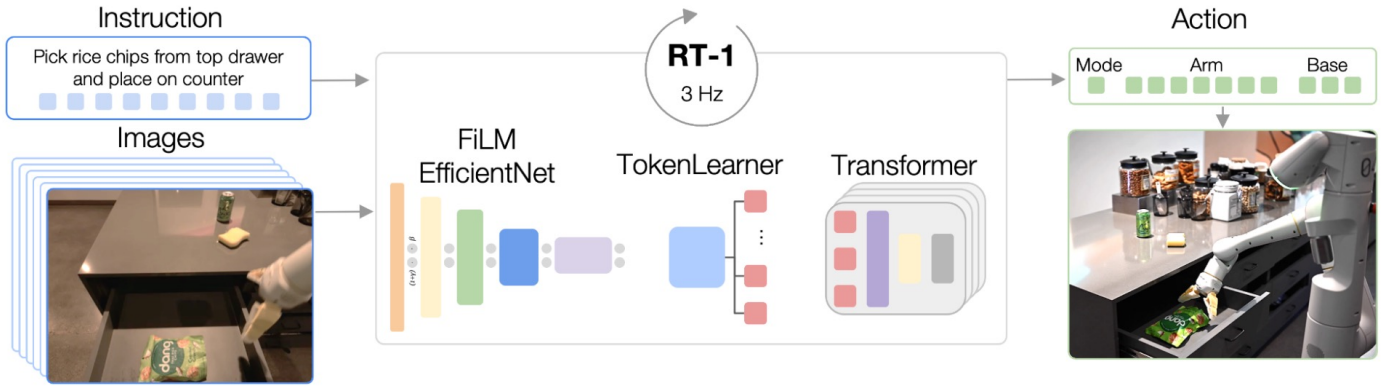

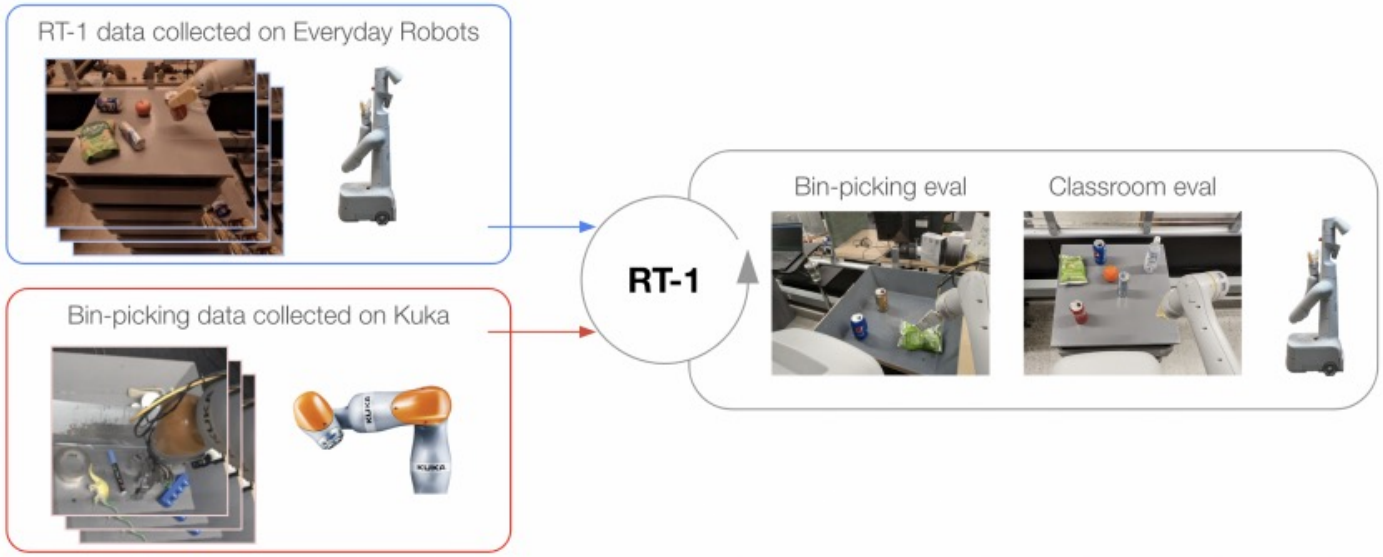

Brohan et al. **RT-1: Robotics Transformer.** 2023.

One policy to control many robots?

the question

Can we create a dataset and model that can generalize in **zero shot** to control entirely new robots?

the setup

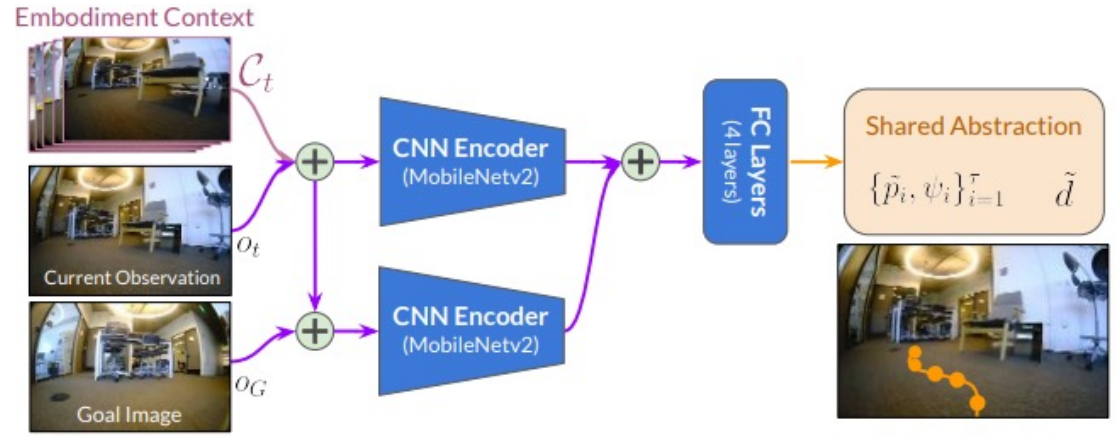

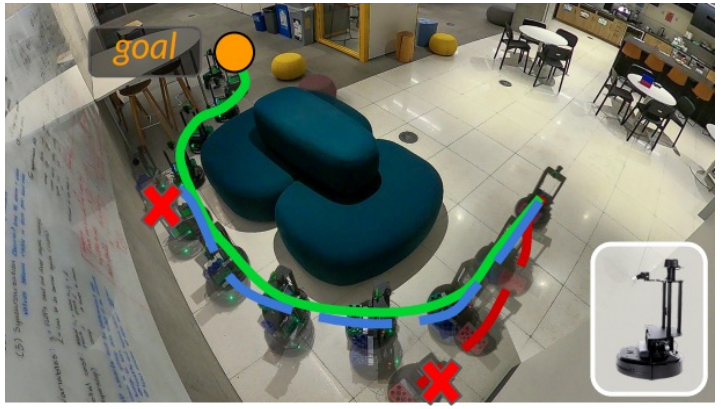

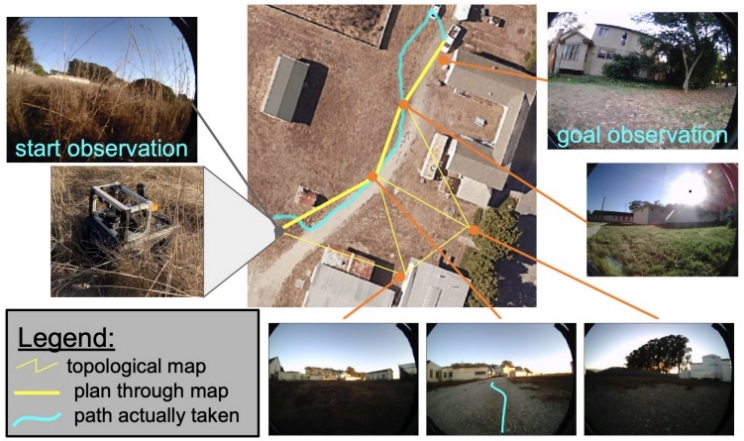

Shah*, Sridhar*, Bhorkar, Hirose, Levine. **GNM: A General Navigation Model to Drive Any Robot**. 2022.

One policy to control many robots?

	Dataset	Platform	Speed	Amt.	Environment
1	GoStanford [26]	TurtleBot2	0.5m/s	14h	office
2	RECON [32]	Jackal	1m/s	25h	off-road
3	CoryHall [35]	RC Car	1.2m/s	2h	hallways
4	Berkeley [33]	Jackal	2m/s	4h	suburban
5	SCAND-S [36]	Spot	1.5m/s	8h	sidewalks
6	SCAND-J [36]	Jackal	2m/s	1h	sidewalks
7	Seattle [37]	Warthog	5m/s	1h	off-road
8	TartanDrive [38]	ATV	10m/s	5h	off-road
	Ours			60h	

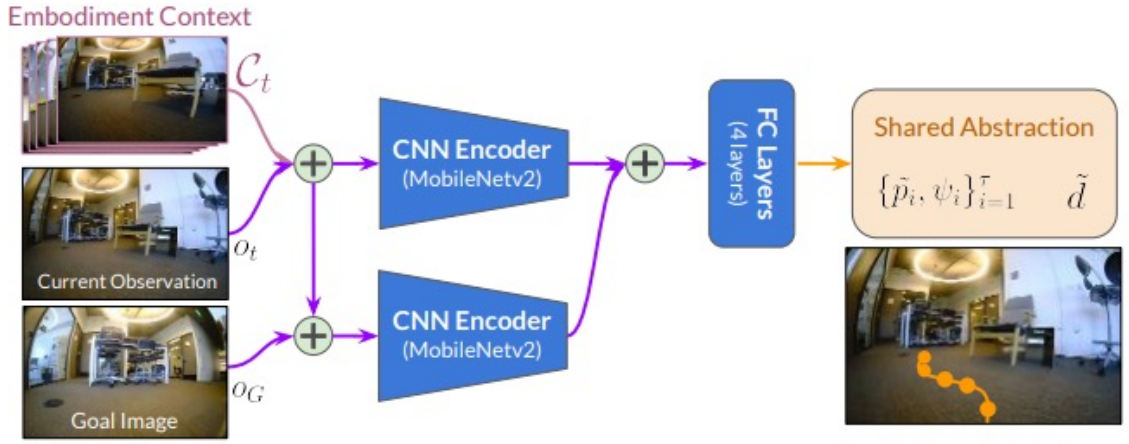

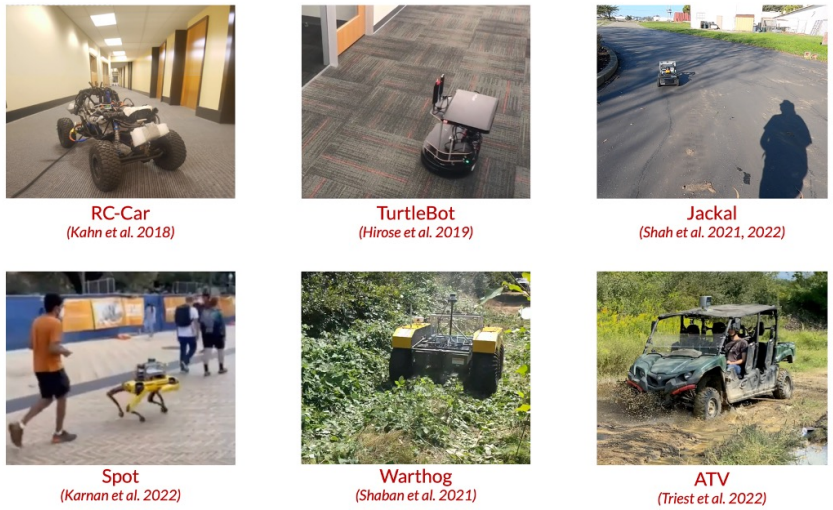

Shah*, Sridhar*, Bhorkar, Hirose, Levine. **GNM: A General Navigation Model to Drive Any Robot**. 2022.

Now make it go fast!

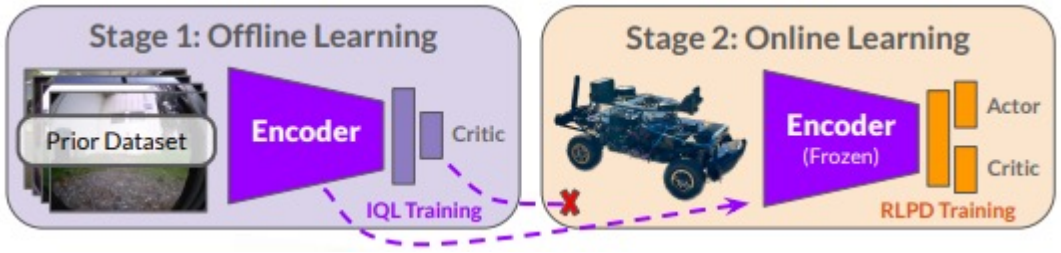

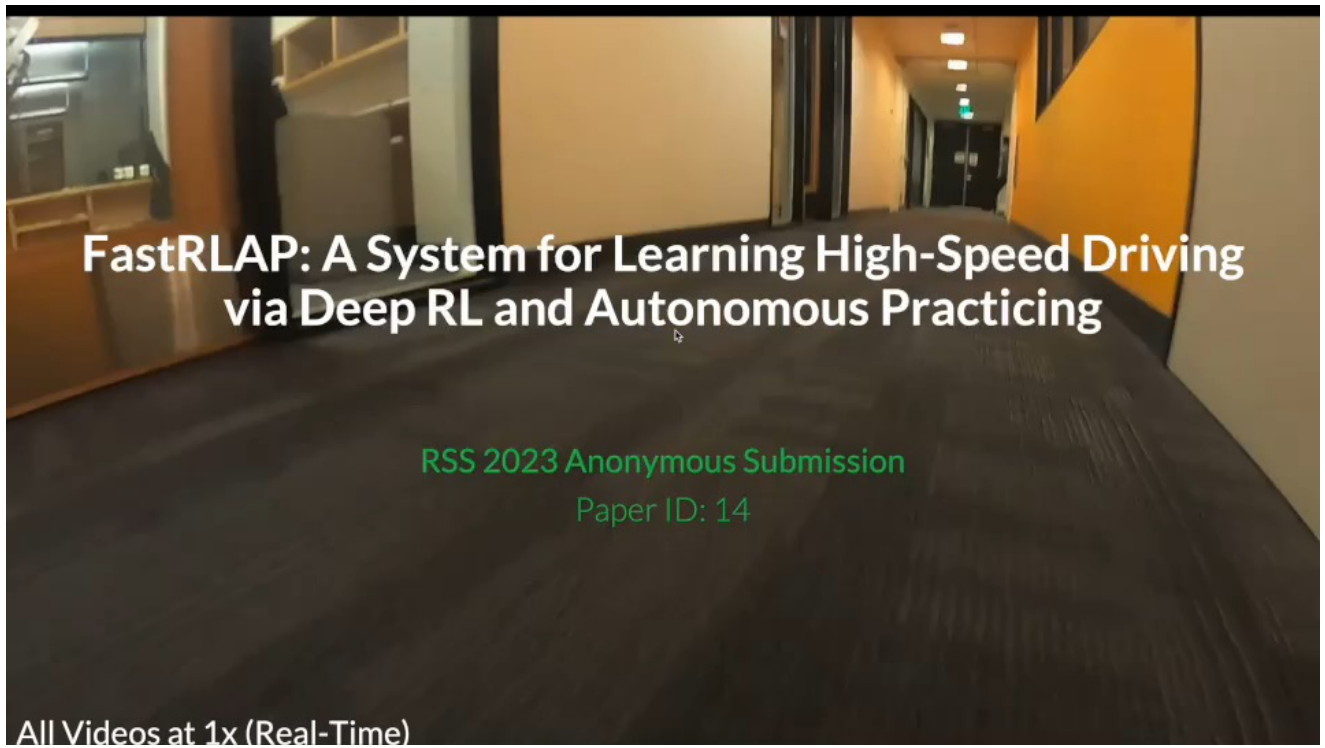

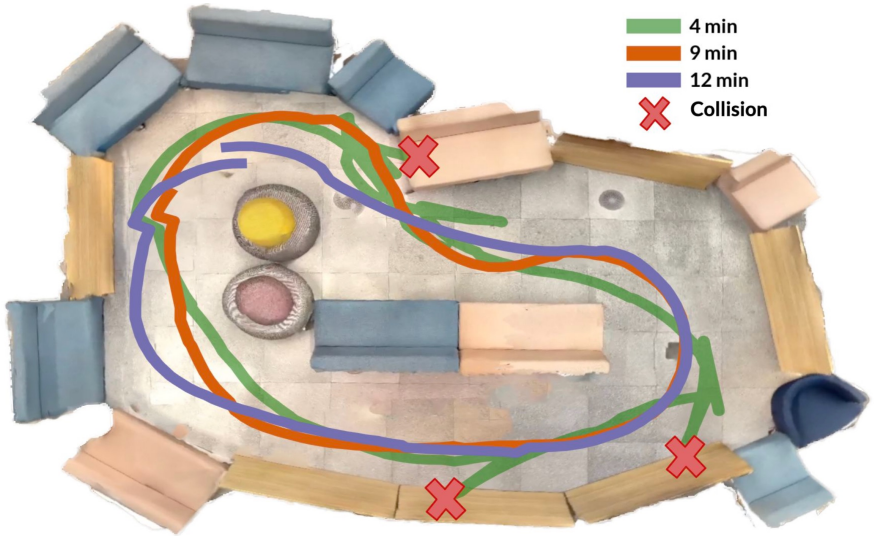

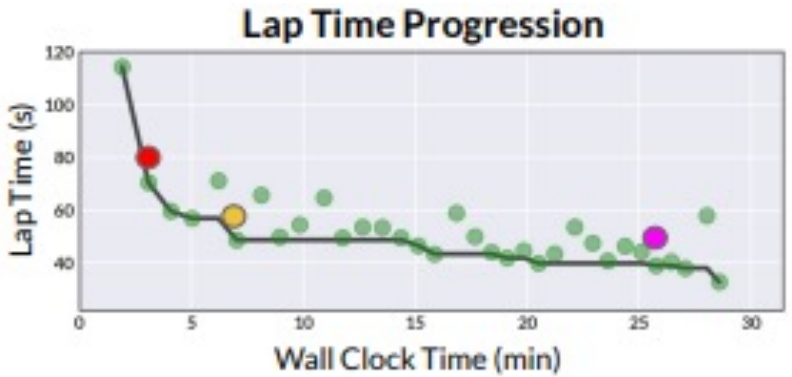

Stachowicz, Bhorkar, Shah, Kostrikov, Levine. **FastRLAP: A System for Learning High-Speed Driving via Deep RL and Autonomous Practicing**. 2023.

How far have we gotten?

✗	Many robots
✓	Many environments
✓	Many tasks
✗	Goal-conditioned
✓	Fine-tunable

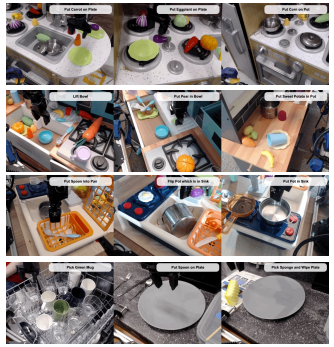

Bridge data & pre-training for robots (PTR)

✗	Many robots
✓	Many environments
✓	Many tasks
✓	Goal-conditioned
✓	Fine-tunable

Bridge data & goal-conditioned RL (FLAP)

✓	Many robots
✓	Many environments
?	Many tasks
✓	Goal-conditioned
✓	Fine-tunable

General navigation models (GNM, FastRLAP)

How robotic learning works now

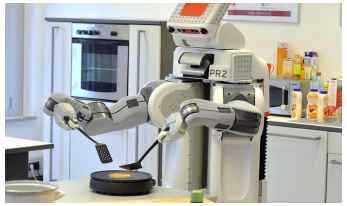

PR2 pancake model

WAM pancake model

UR10 box picking model

lipstick robot model

mousetrap… hand… shake… model??

How robotic learning will work in the future

giant un/self-supervised pretrained **robot** model

"robot foundation model"

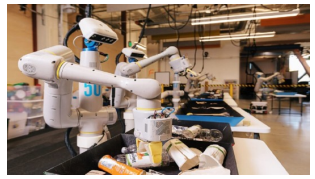

finetuning

trash sorting model

What would it take to build such a general-purpose robot model?

RAIL
Robotic AI & Learning Lab

website: http://rail.eecs.berkeley.edu
source code: http://rail.eecs.berkeley.edu/code.html