

Data for Robot Manipulation: From Demonstrations to Autonomy

Chelsea Finn



How can robots **generalize** to the real world?

How can robots **generalize** to the real world?

Recipe: Large, expressive models + Broad data

Key challenge

Where does the data come from?

Demonstration data? Autonomously collected data?

How we get broad robotics data?

What robot learning research currently does

Collect a new data for the project.

Don't use the data again.

for *each* and *every* project

What ML research currently does

Collect a large dataset **once**
(e.g. ImageNet, Wikitext)

Reuse it many times

Or even better: Reuse a pre-trained model

Can we share *data* and *pre-trained models* **across institutions** in robotics?

If so, I think we could make a lot more progress collectively!

What is stopping us from sharing data across institutions?

your next experiment

RoboTurk dataset

MIME dataset

Robonet dataset

Bridge dataset

the experiment after that



What is stopping us from sharing data across institutions?

Can we make a **larger island** in a **centralized, community-driven effort**?


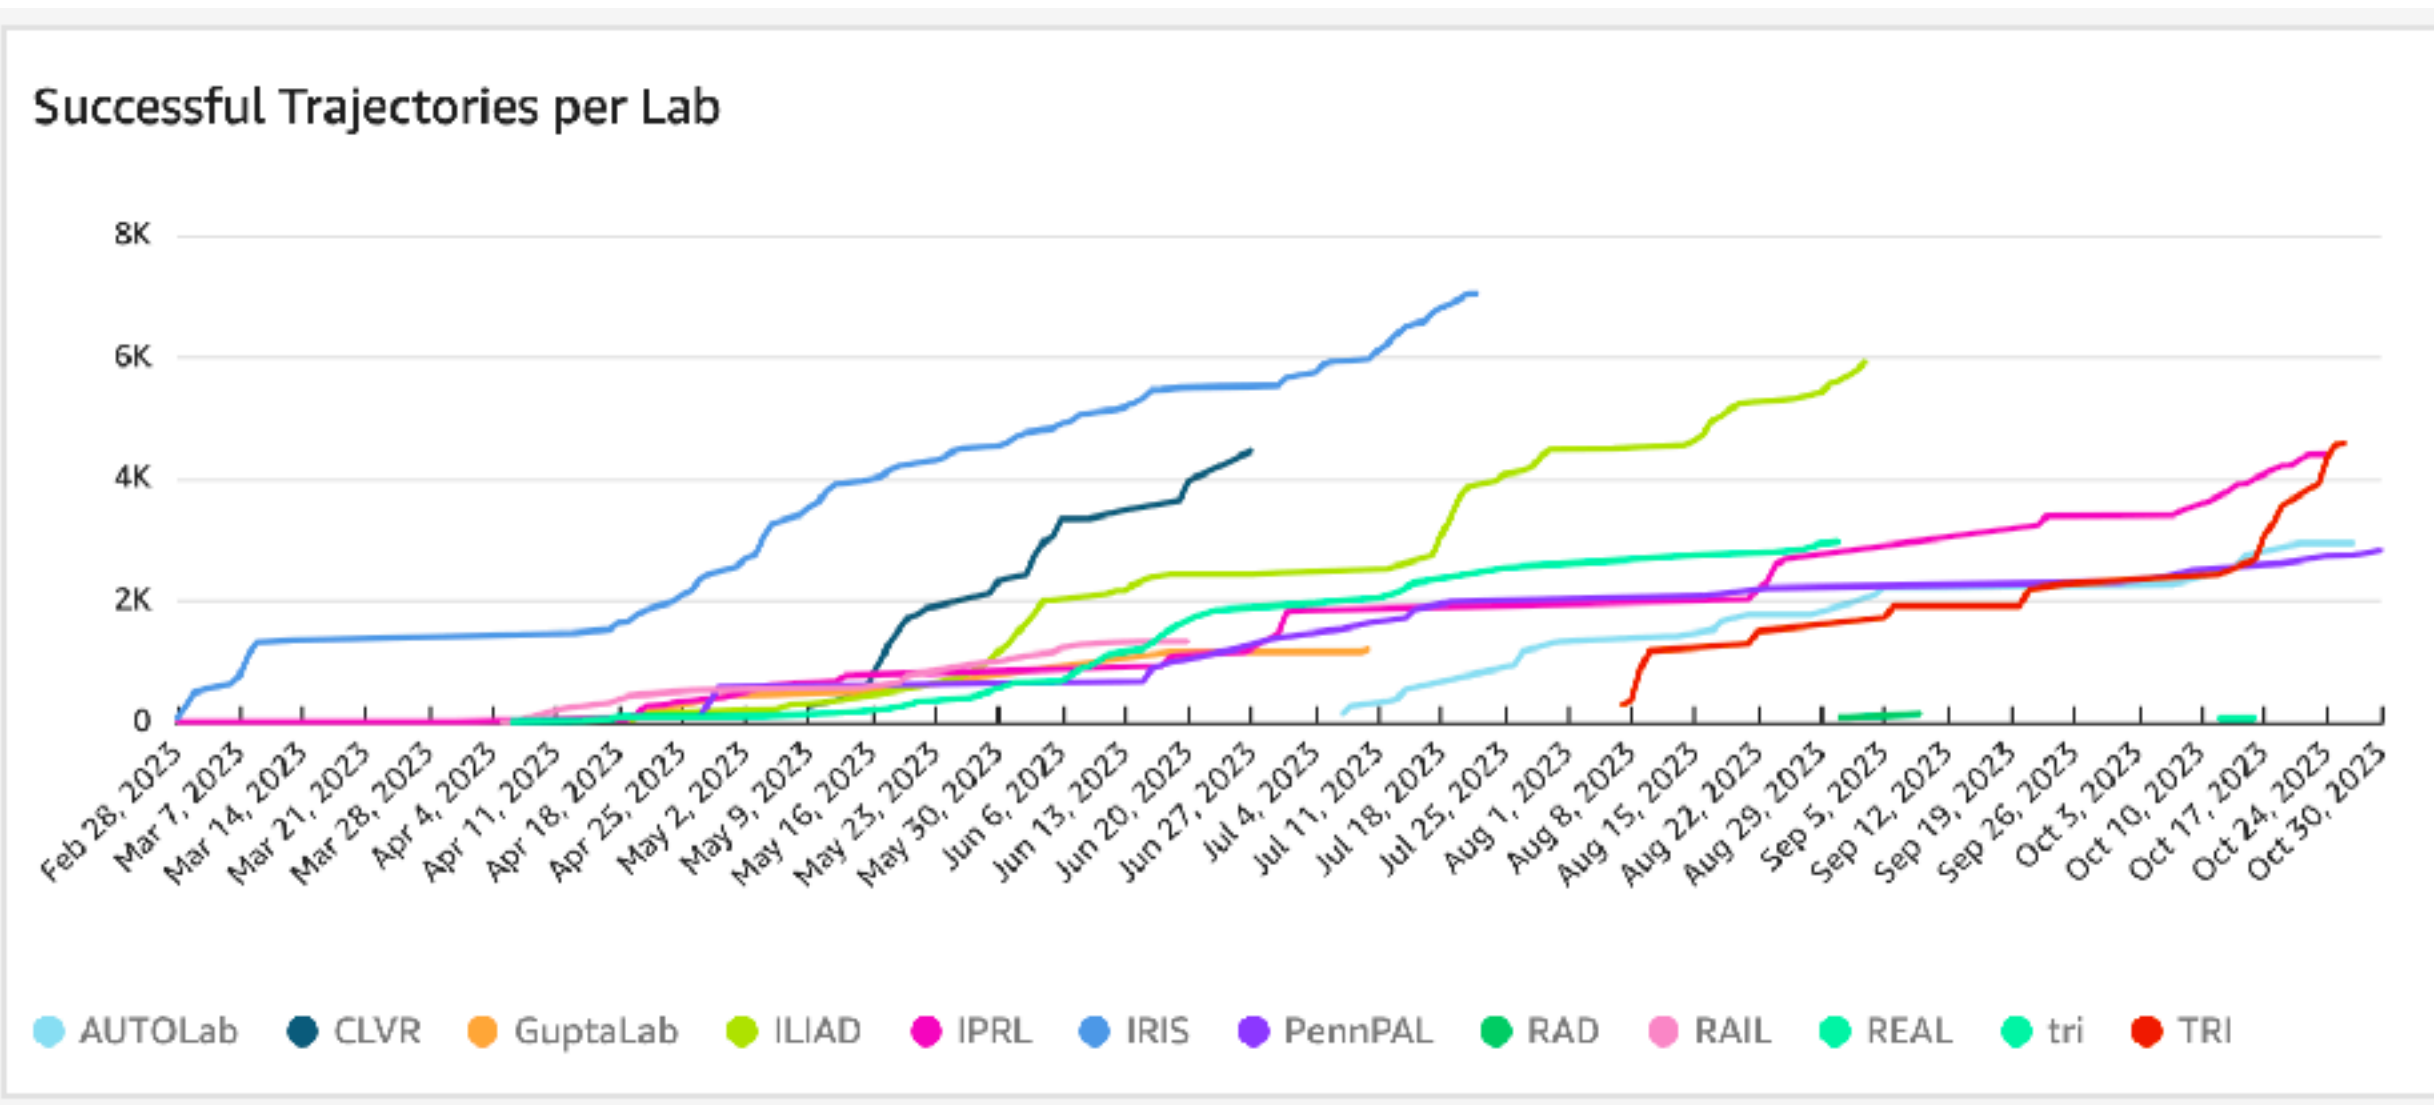
Project Slack: 70 people from 18 slack orgs

What is stopping us from sharing data across institutions?

Can we make a larger island in a centralized, community-driven effort?


Project Slack: 70 people from 18 slack orgs

So far: 38k+ demos in 600+ scenes
collected by 11 labs on common platform



What is stopping us from sharing data across institutions?

Let's also use the data we already have!




Step 1: Put all of the data in a common format

Step 2: Train a large model on the data

What is stopping us from sharing data across institutions?

Let's also use the data we already have!





Step 1: Put all of the data in a common format ...as a community!

As of Oct 8:

- 174 co-authors
- 34 research labs
- 21 institutions



The Aggregated Open X-Embodiment Dataset



1M+ Real Robot Episodes



22 Robot Embodiments



34 Research Labs



300+ Scenes

Using the Open X-Embodiment Data

- Download & sampling is straightforward:

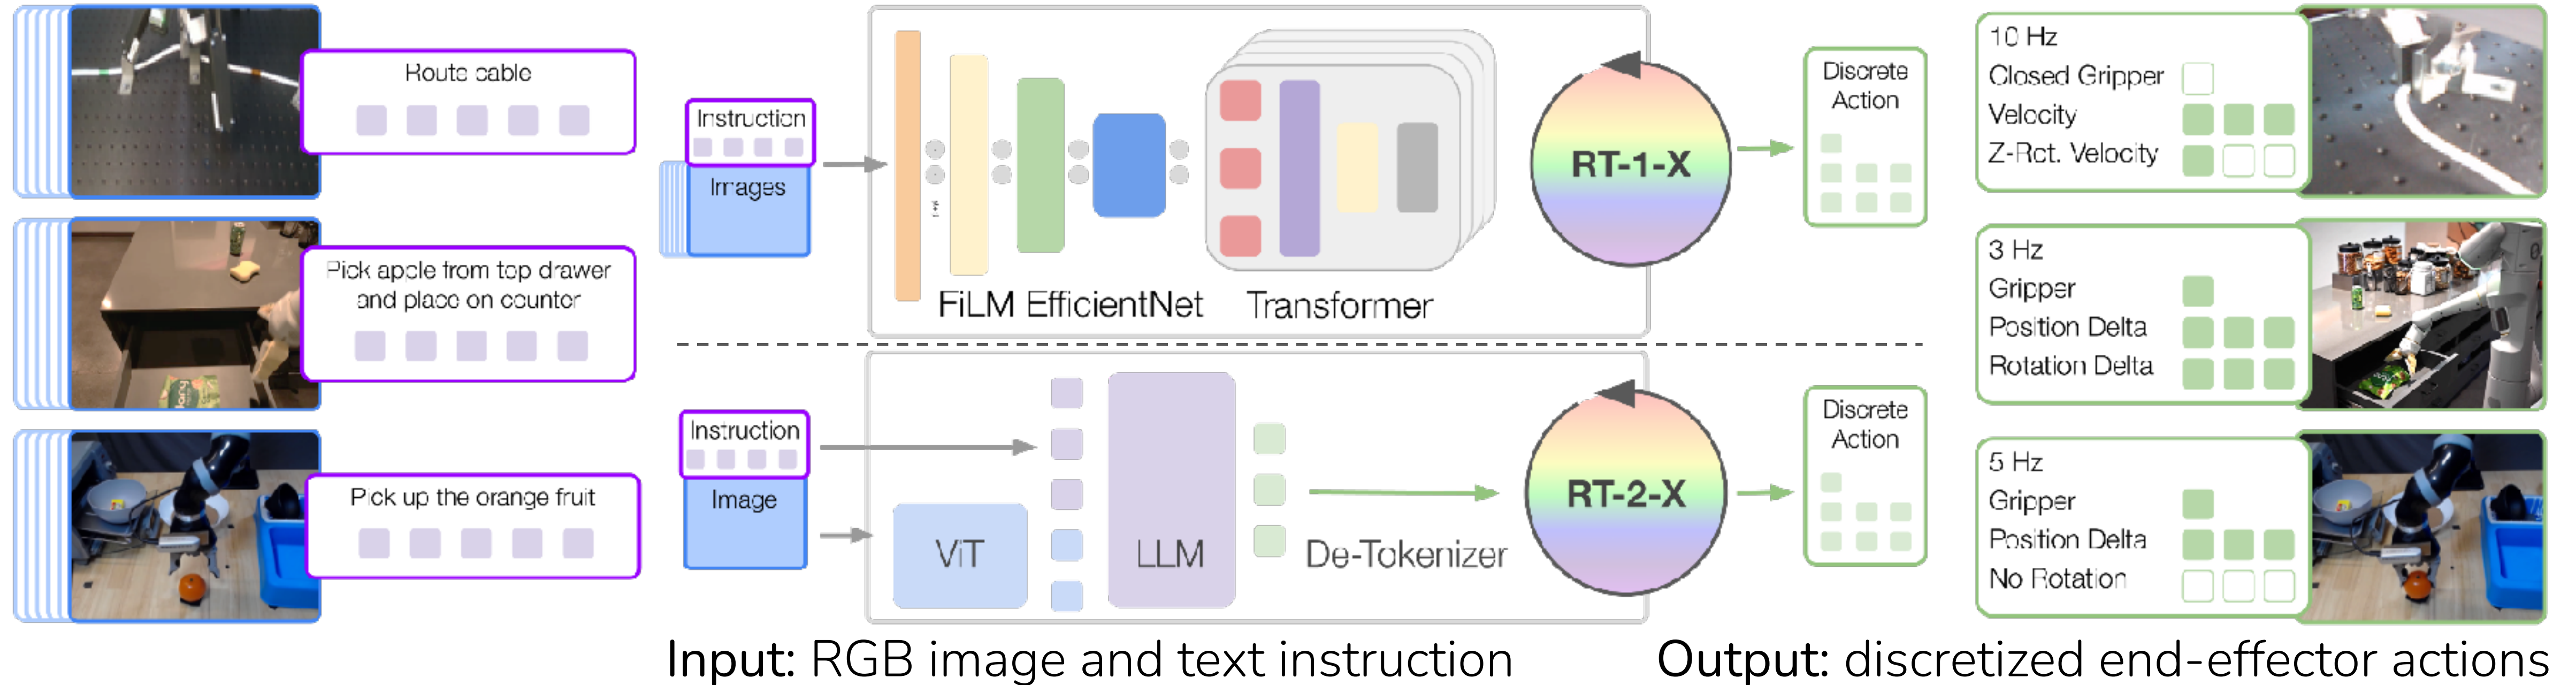
```
▶ import tensorflow_datasets as tfds

ds = tfds.load('bridge', split='train')
for episode in ds:
    for step in episode['steps']:
        image = step['observation']['image_0']
        action = step['action']
        ...
```

For details, download, & colabs, see: robotics-transformer-x.github.io

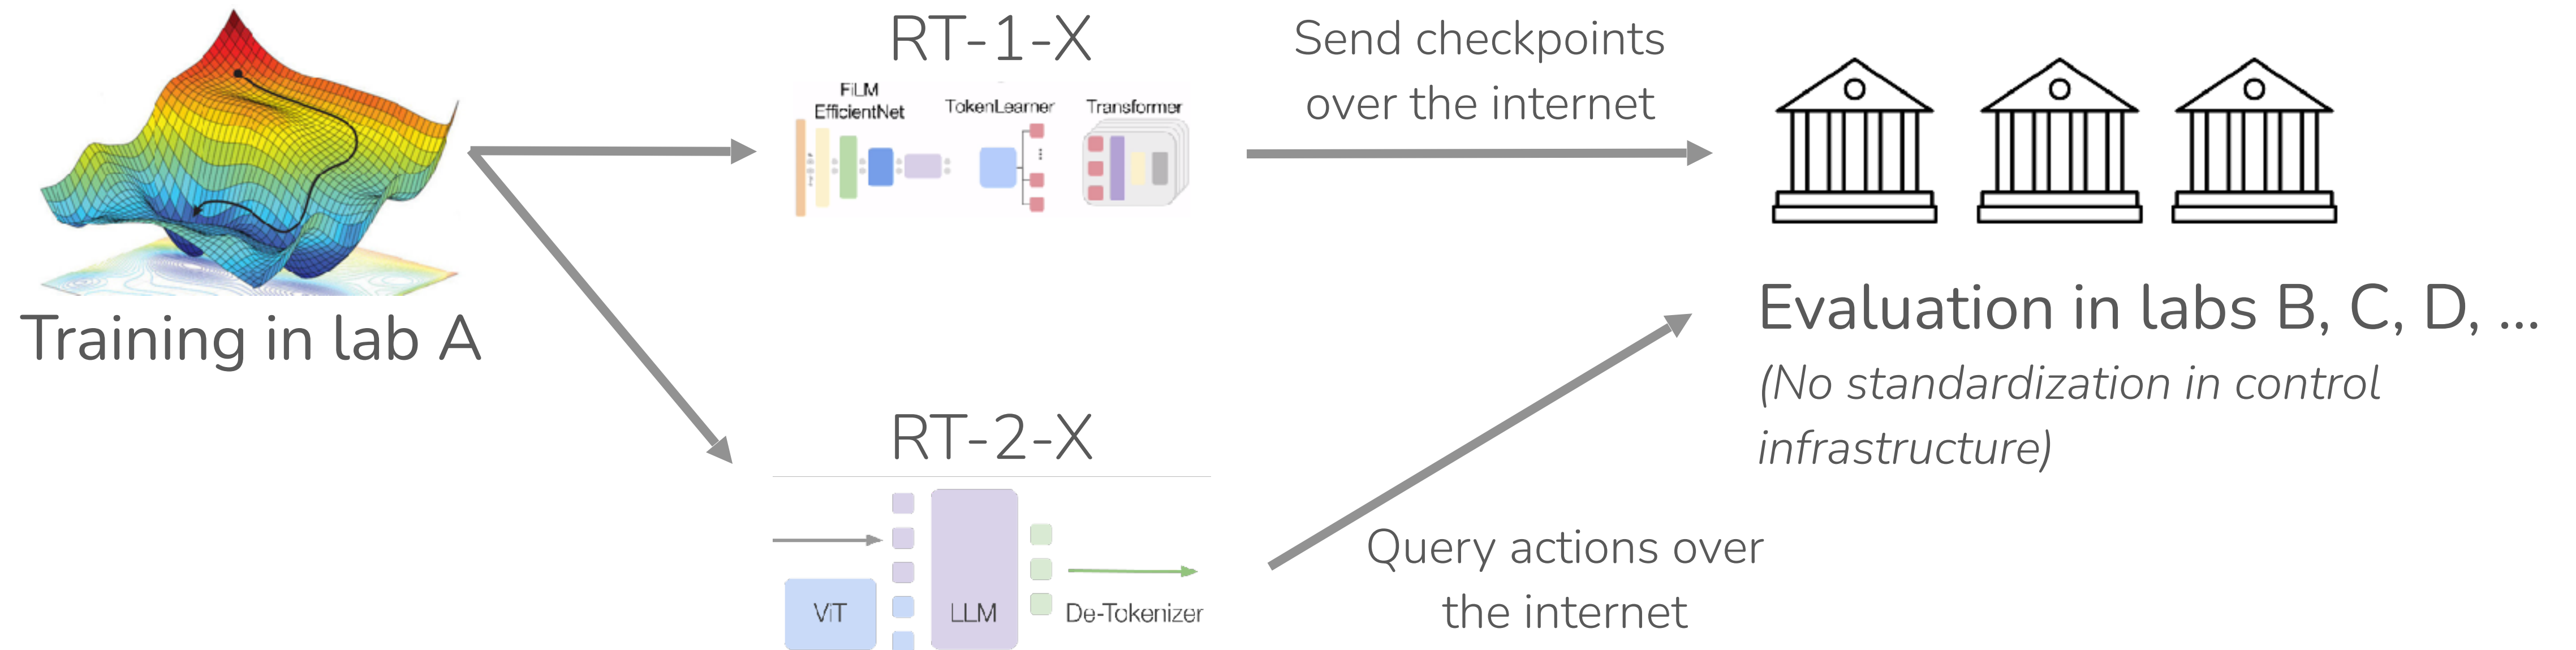
Can we train models on the data?

Reuse architectures from prior works.




- Trained using supervised learning.
- RT-2-X uses pre-trained vision-language model

Training & Evaluation Methodology



RT-1-X Experimental Results




Key takeaways

- Training on cross-embodiment data improves over training on in-distribution data.
- Including the original method from each lab.

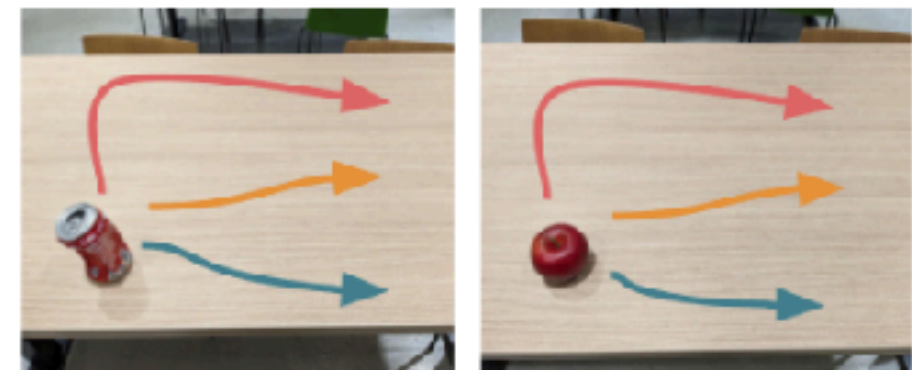
20% average improvement

RT-2-X Results: Novel Instructions Outside of Google Robot Data


(a) Absolute Motion
move the chip bag to the **top** / **bottom** right of the counter




move to **top right** / **right** / **bottom right**



(b) Object-Relative Motion
move apple between **coke and cup** / **coke and sponge** / **cup and sponge**



(c) Preposition Alters Behavior
put apple **on** cloth /
move apple **near** cloth




put orange **into** the pot /
move orange **near** pot




put banana **on top of** the pan /
move banana **near** pan



move apple between can & orange



■ RT-2 ■ RT-2-X



Takeaway: Large-scale data is starting to become a reality for robotic manipulation.

Also lots of future opportunity:

Hope for broader community participation,
continued dataset expansion

If you're interested in joining: email open-x-embodiment@googlegroups.com

Come to Thursday 11 am sponsor talk for info on access to RT-2-X.



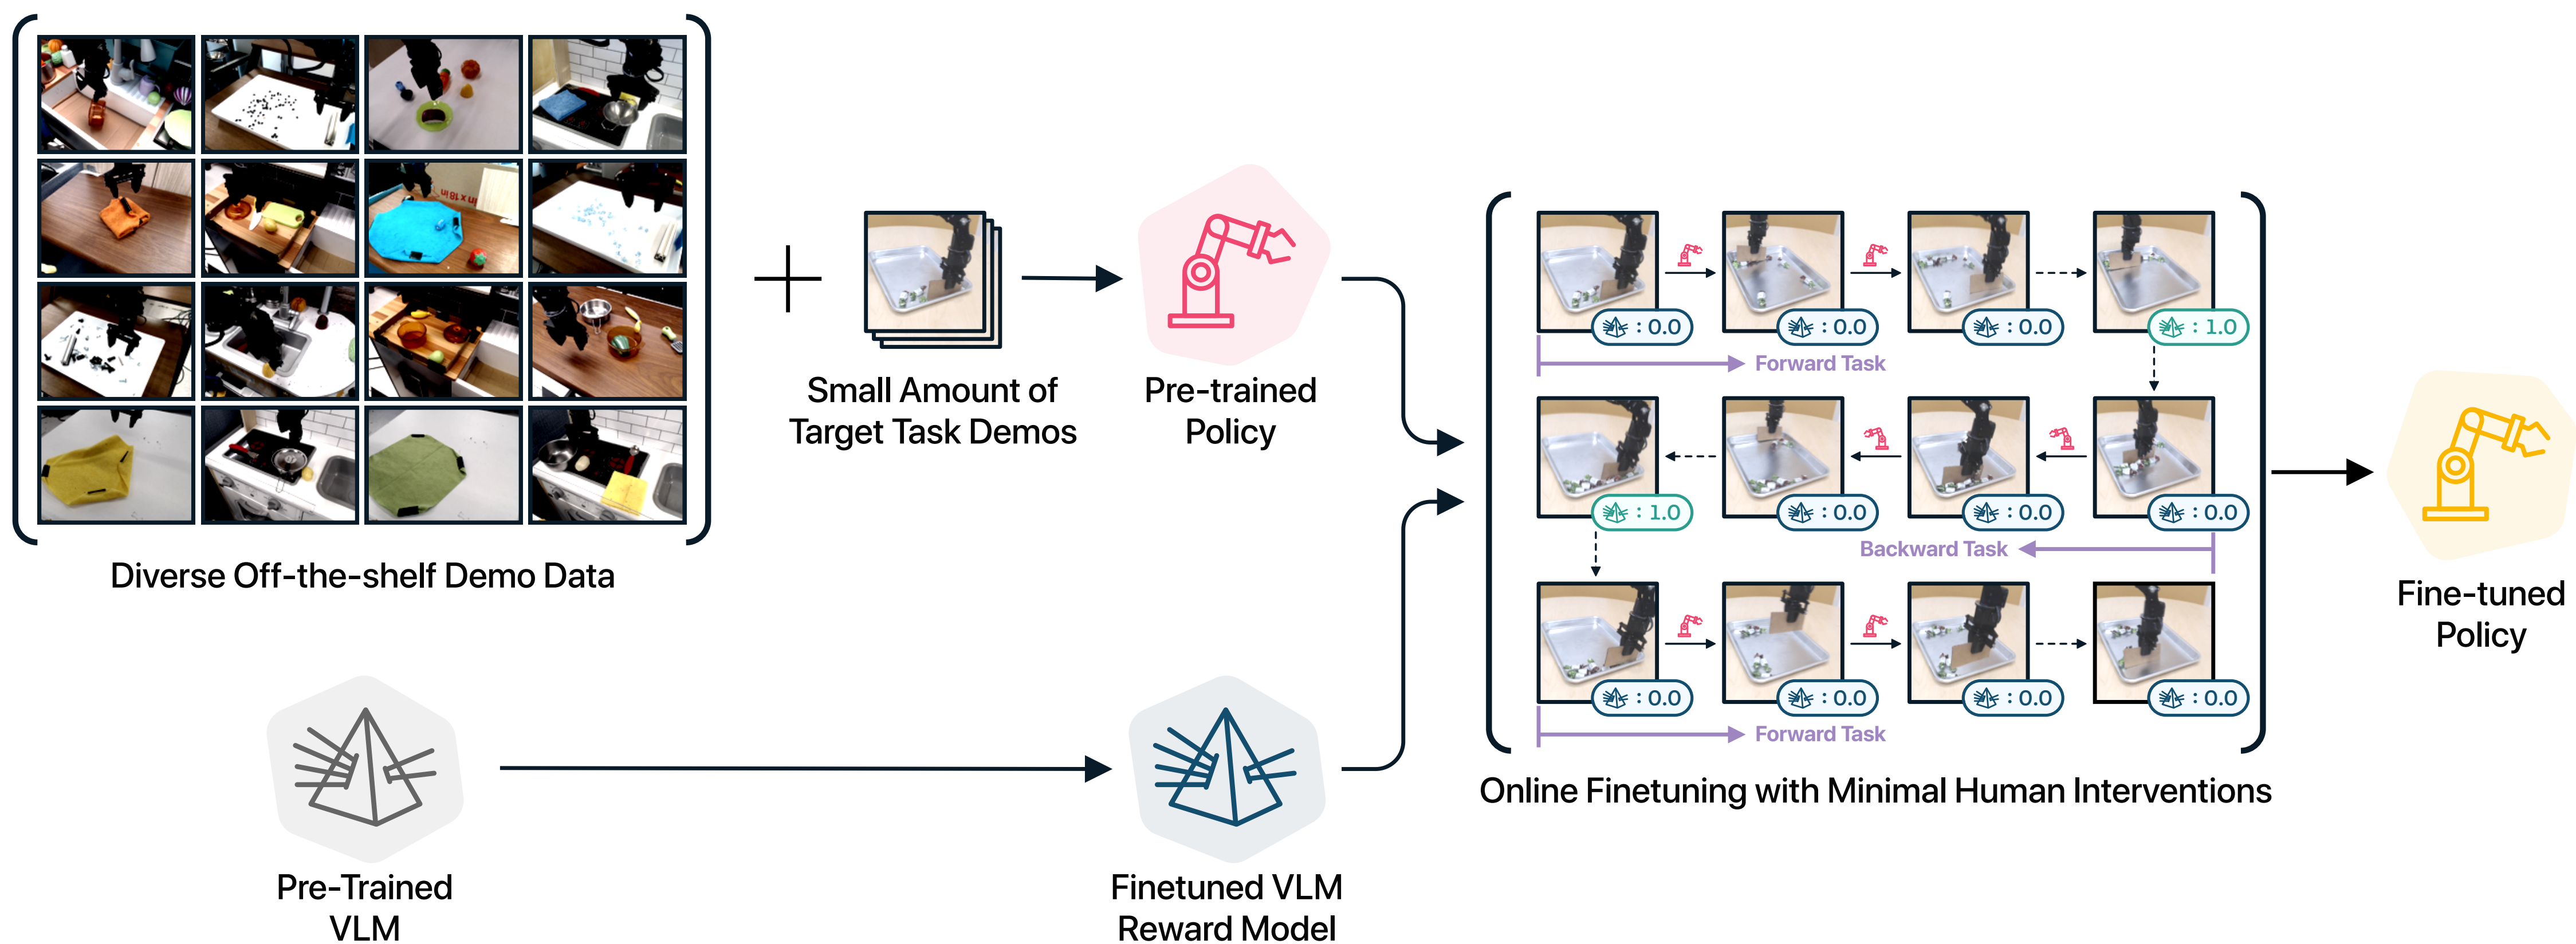
Everything so far: human-teleoperated demonstration data

Everything so far: **human**-teleoperated demonstration data

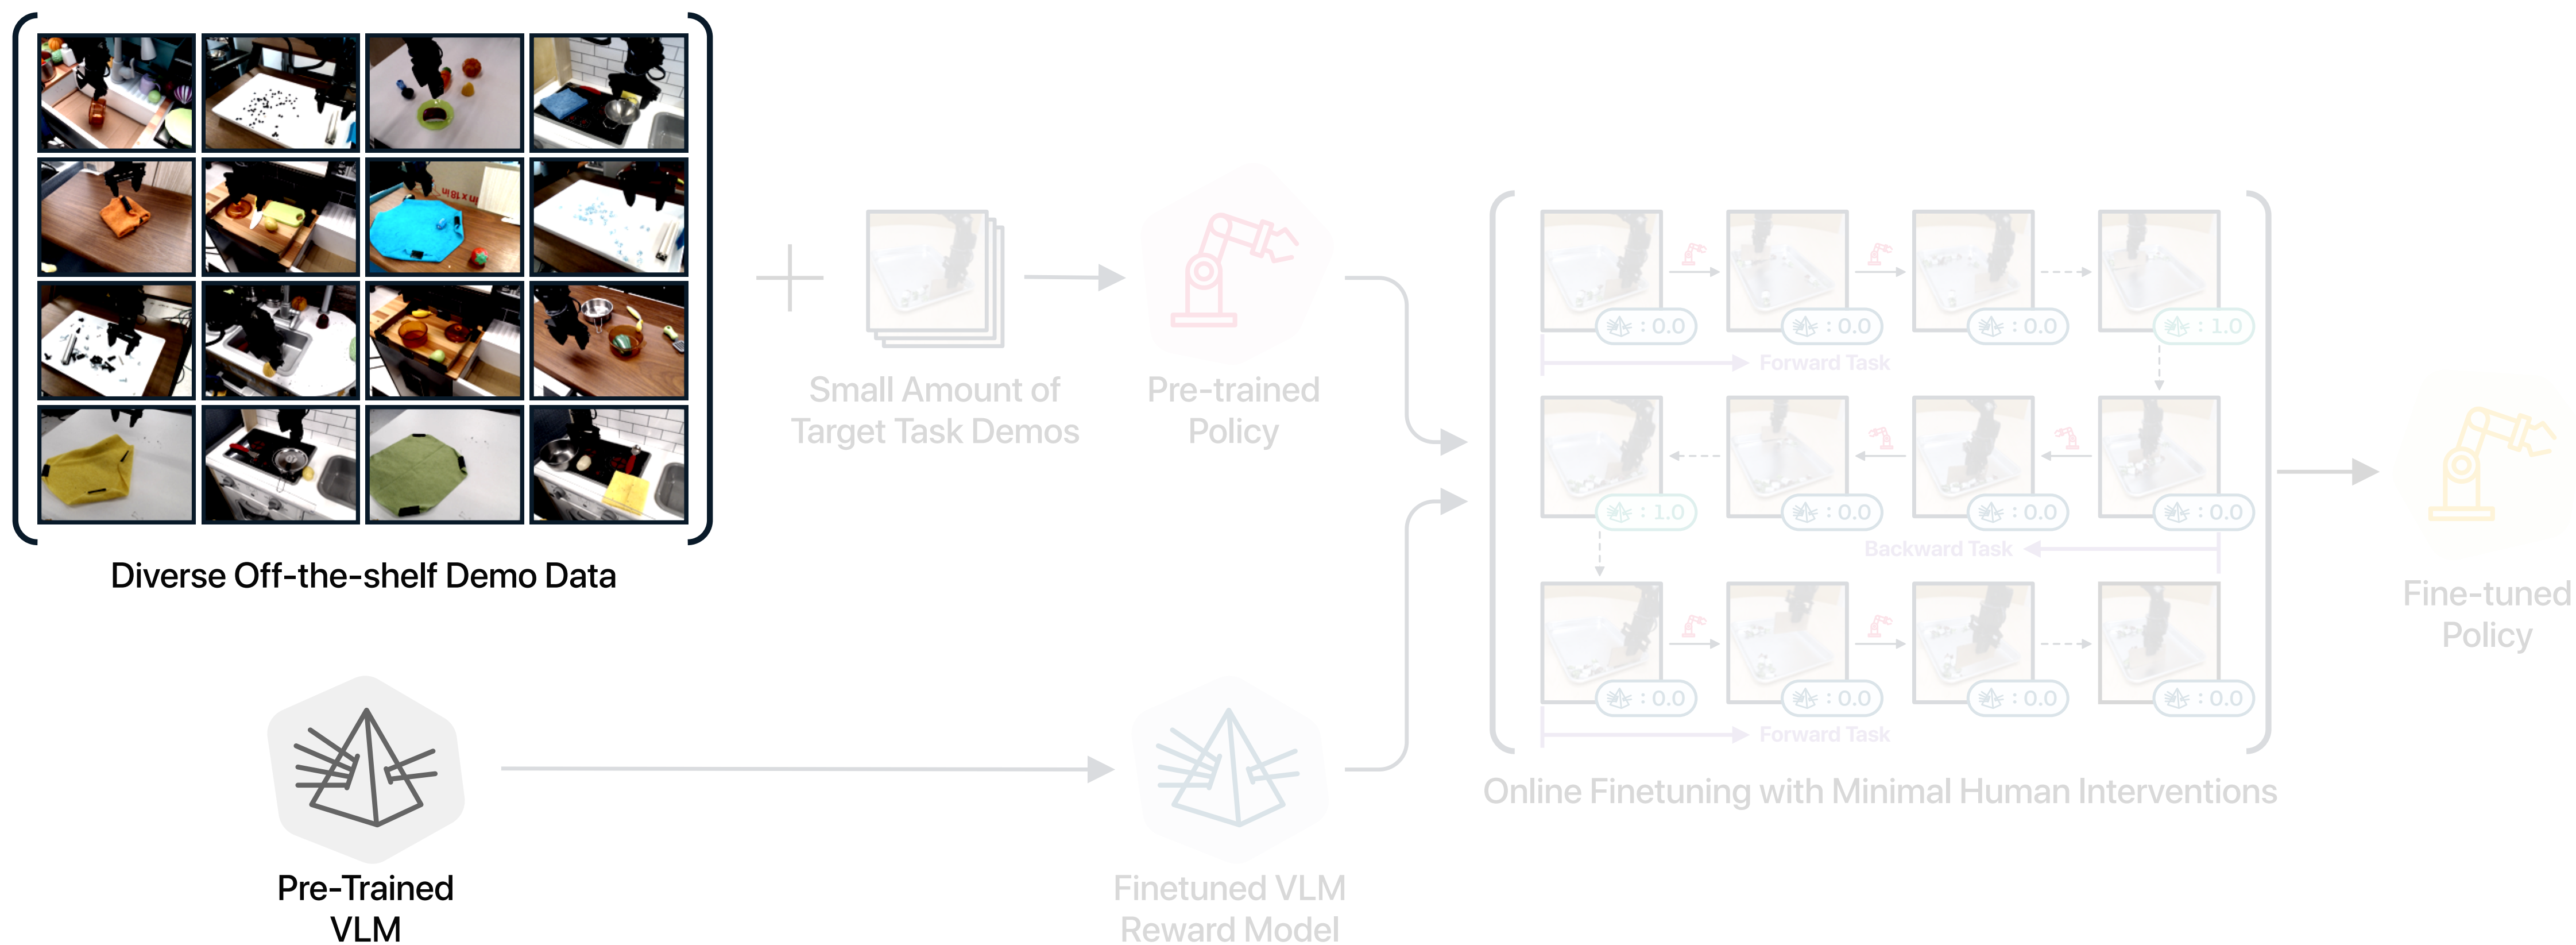
 not scalable

Can we grow the data with **autonomous collection**,
to cover more and more tasks?

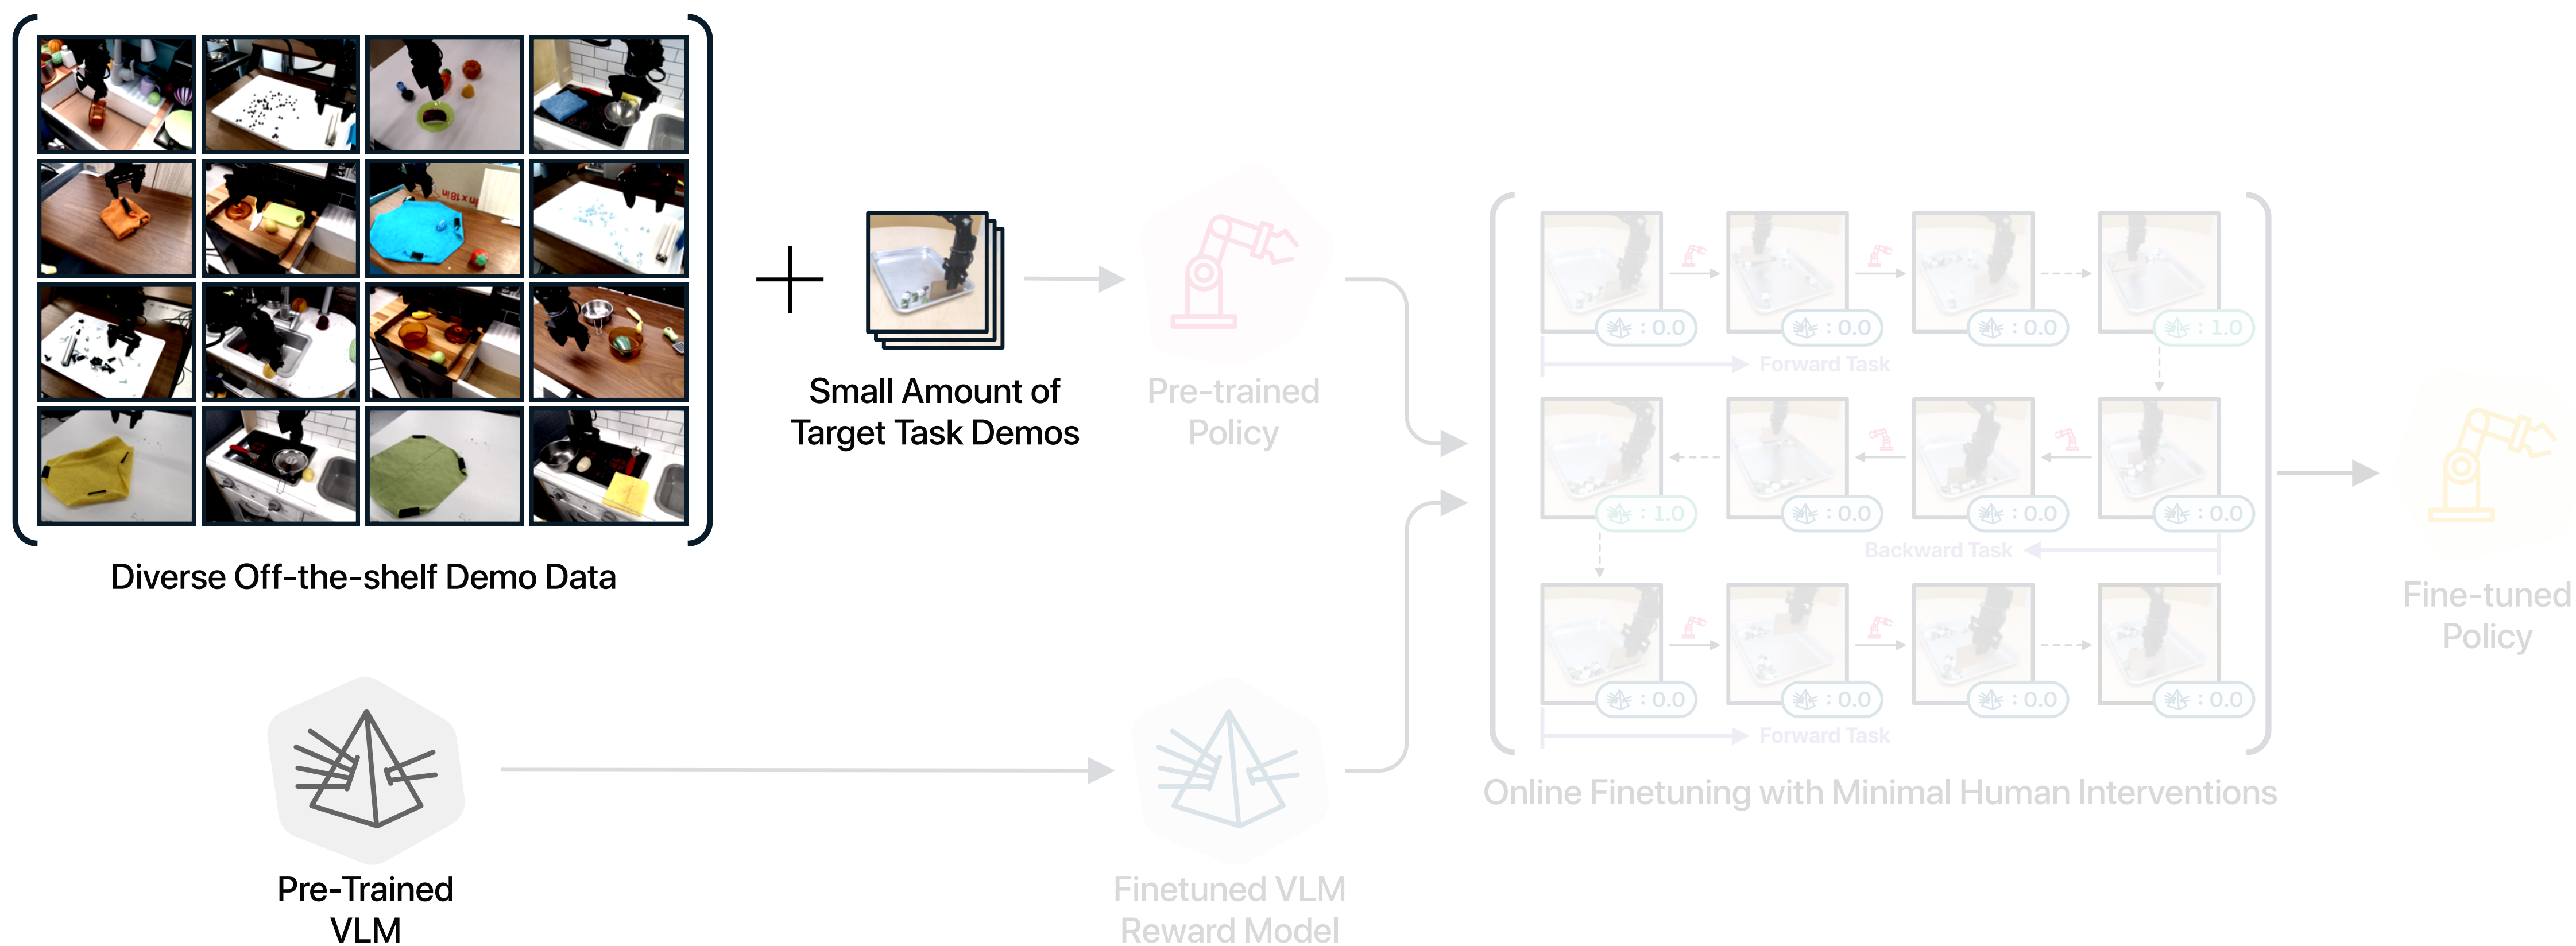
RoboFuME: Making Robot Fine-tuning Easy by Improving Autonomy



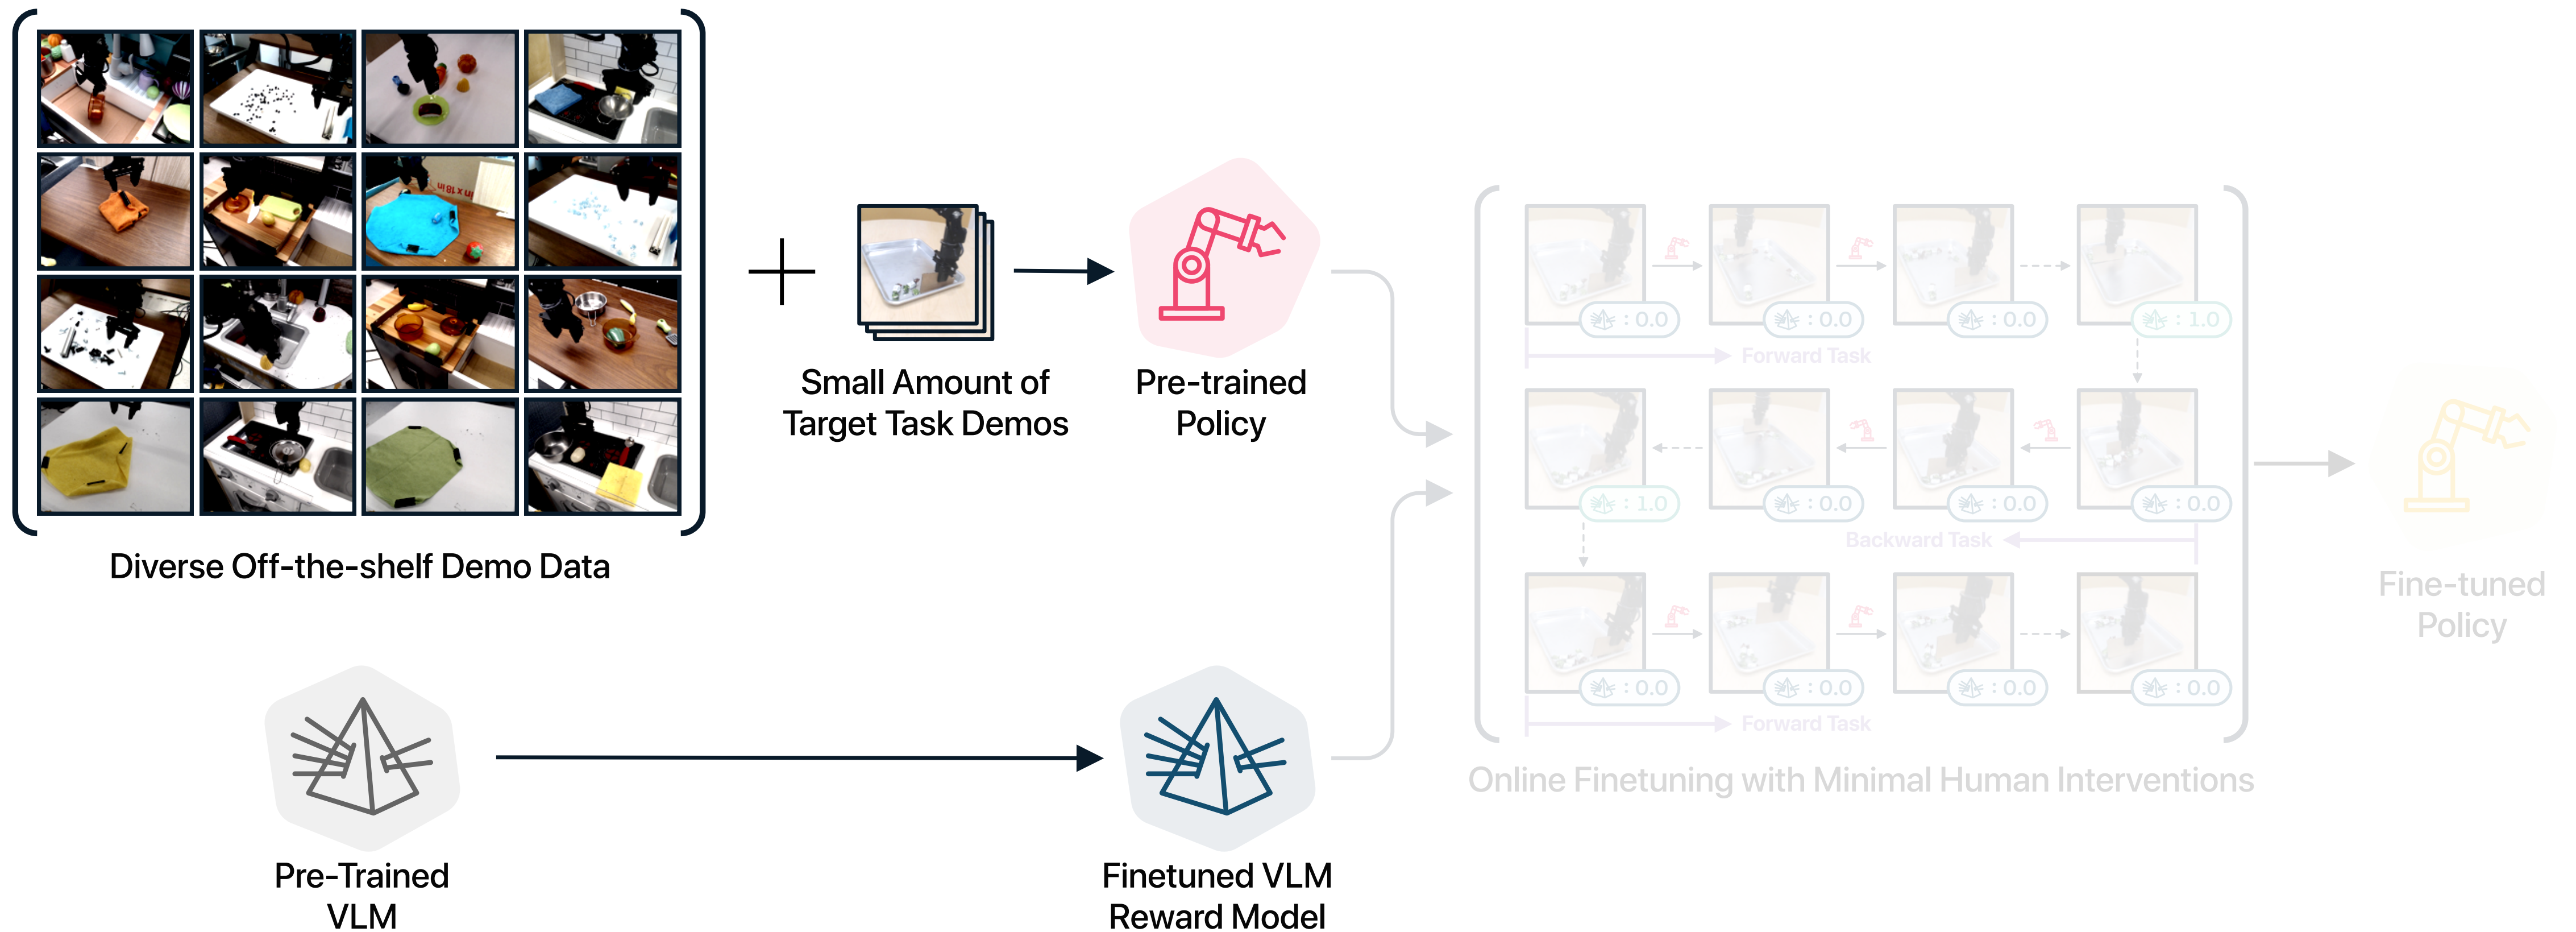
RoboFuME: Making Robot Fine-tuning Easy by Improving Autonomy



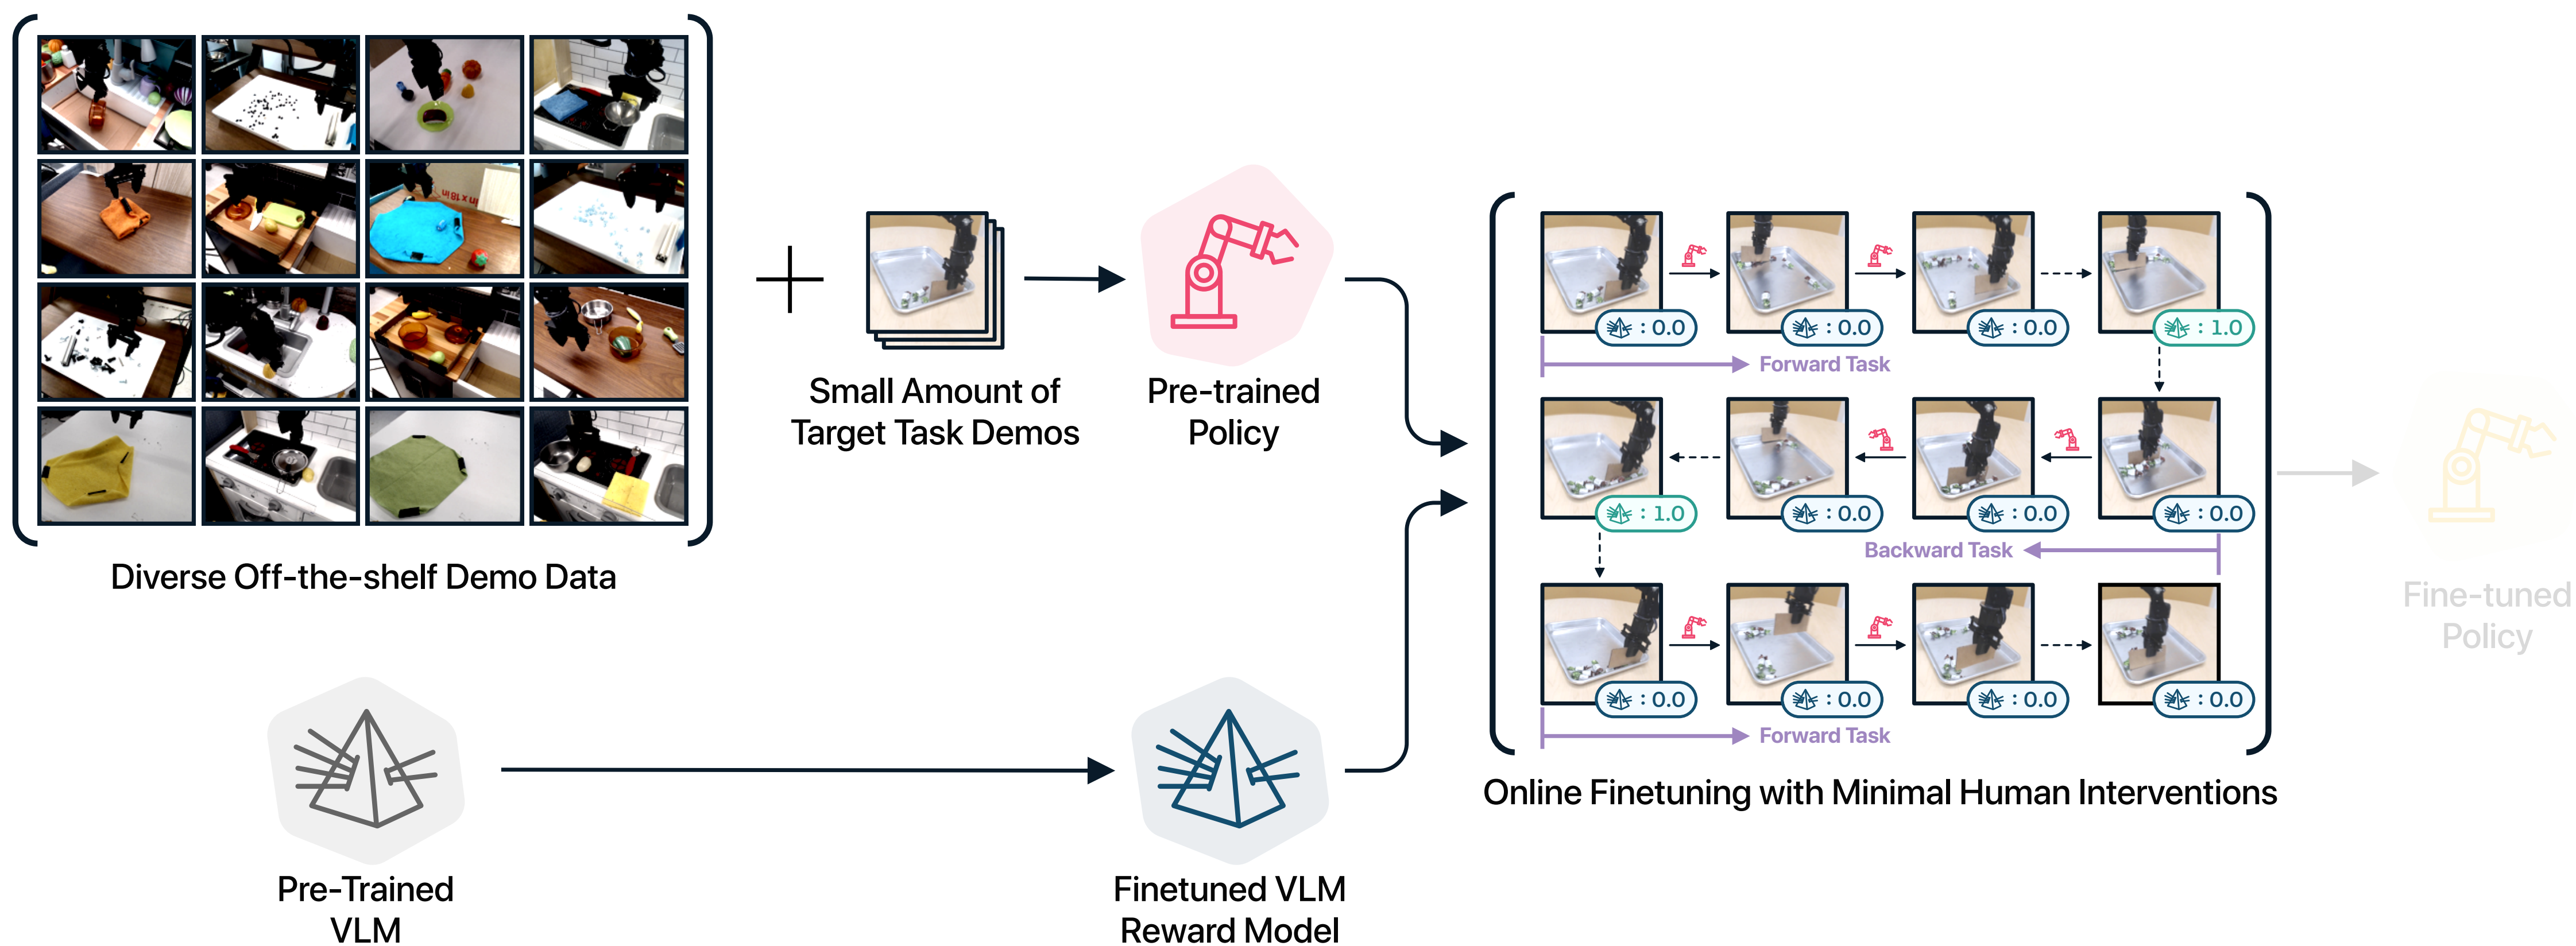
RoboFuME: Making Robot Fine-tuning Easy by Improving Autonomy



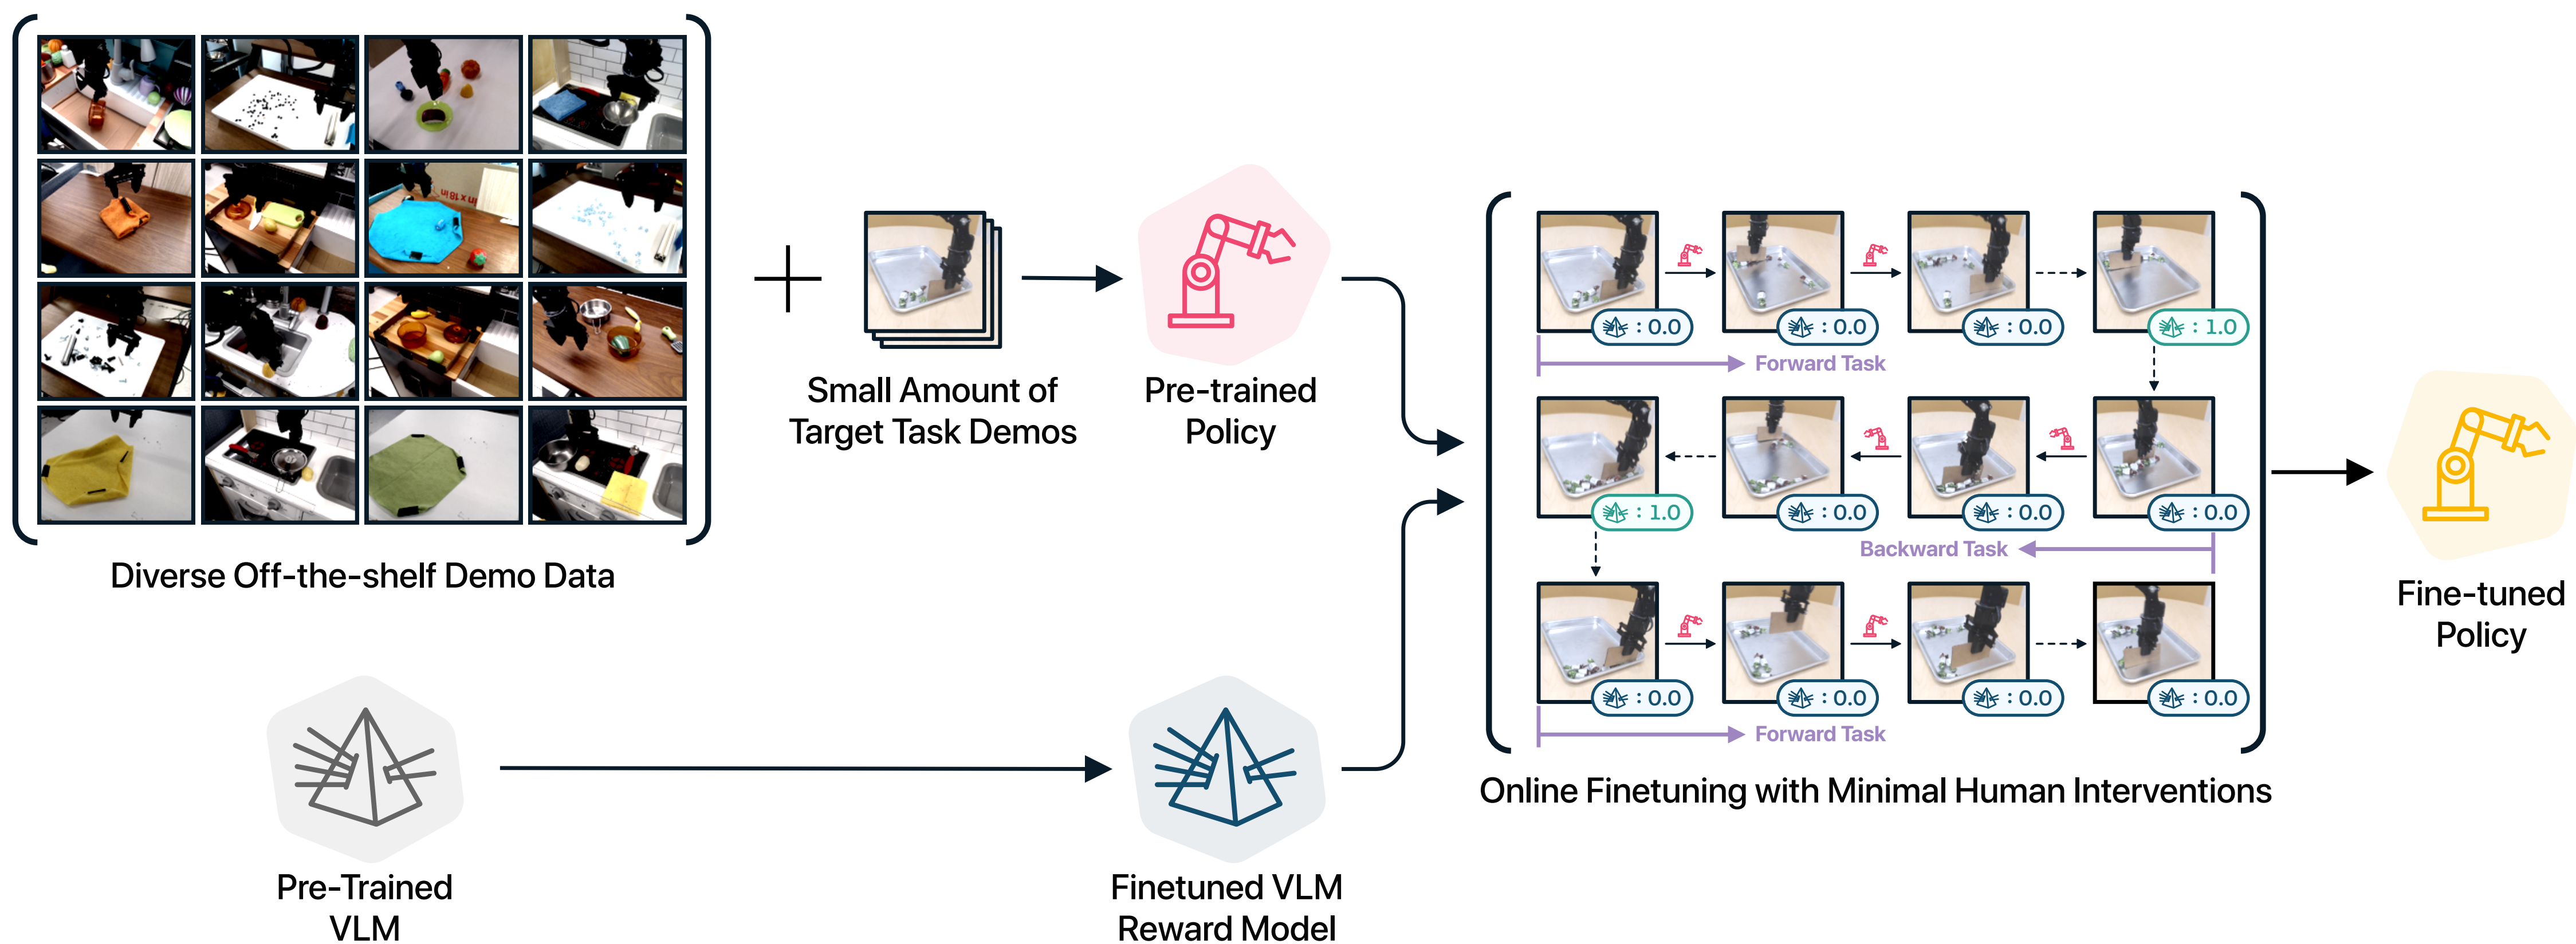
RoboFuME: Making Robot Fine-tuning Easy by Improving Autonomy



RoboFuME: Making Robot Fine-tuning Easy by Improving Autonomy



RoboFuME: Making Robot Fine-tuning Easy by Improving Autonomy



RoboFuME: Key Design Choices

- Learn language-conditioned policy & critic on all data
- Fine-tune MiniGPT4 to predict binary success/failure using prior & new demo data (incl. small # of failed demos)
- [Cal-QL algorithm](#) for RL pre-training and fine-tuning
- Behavior cloning regularization

Real Robot Experiments

Prior Data: Bridge V2 (off-the-shelf, collected at UC Berkeley)



Walke, Homer, et al. "BridgeData V2: A Dataset for Robot Learning at Scale."

Downstream Tasks (to be learned at Stanford)



Candy Sweeping



Cloth Folding



Object Covering



Pot Lid




Pot Pick-and-Place

Demos for Target Tasks (Candy Sweeping)




Successful Forward Demos (Samples)



Successful Backward Demos (Samples)




Failure Forward Demos (Samples)




Failure Backward Demos (Samples)

Reset-free Fine-tuning (Candy Sweeping)

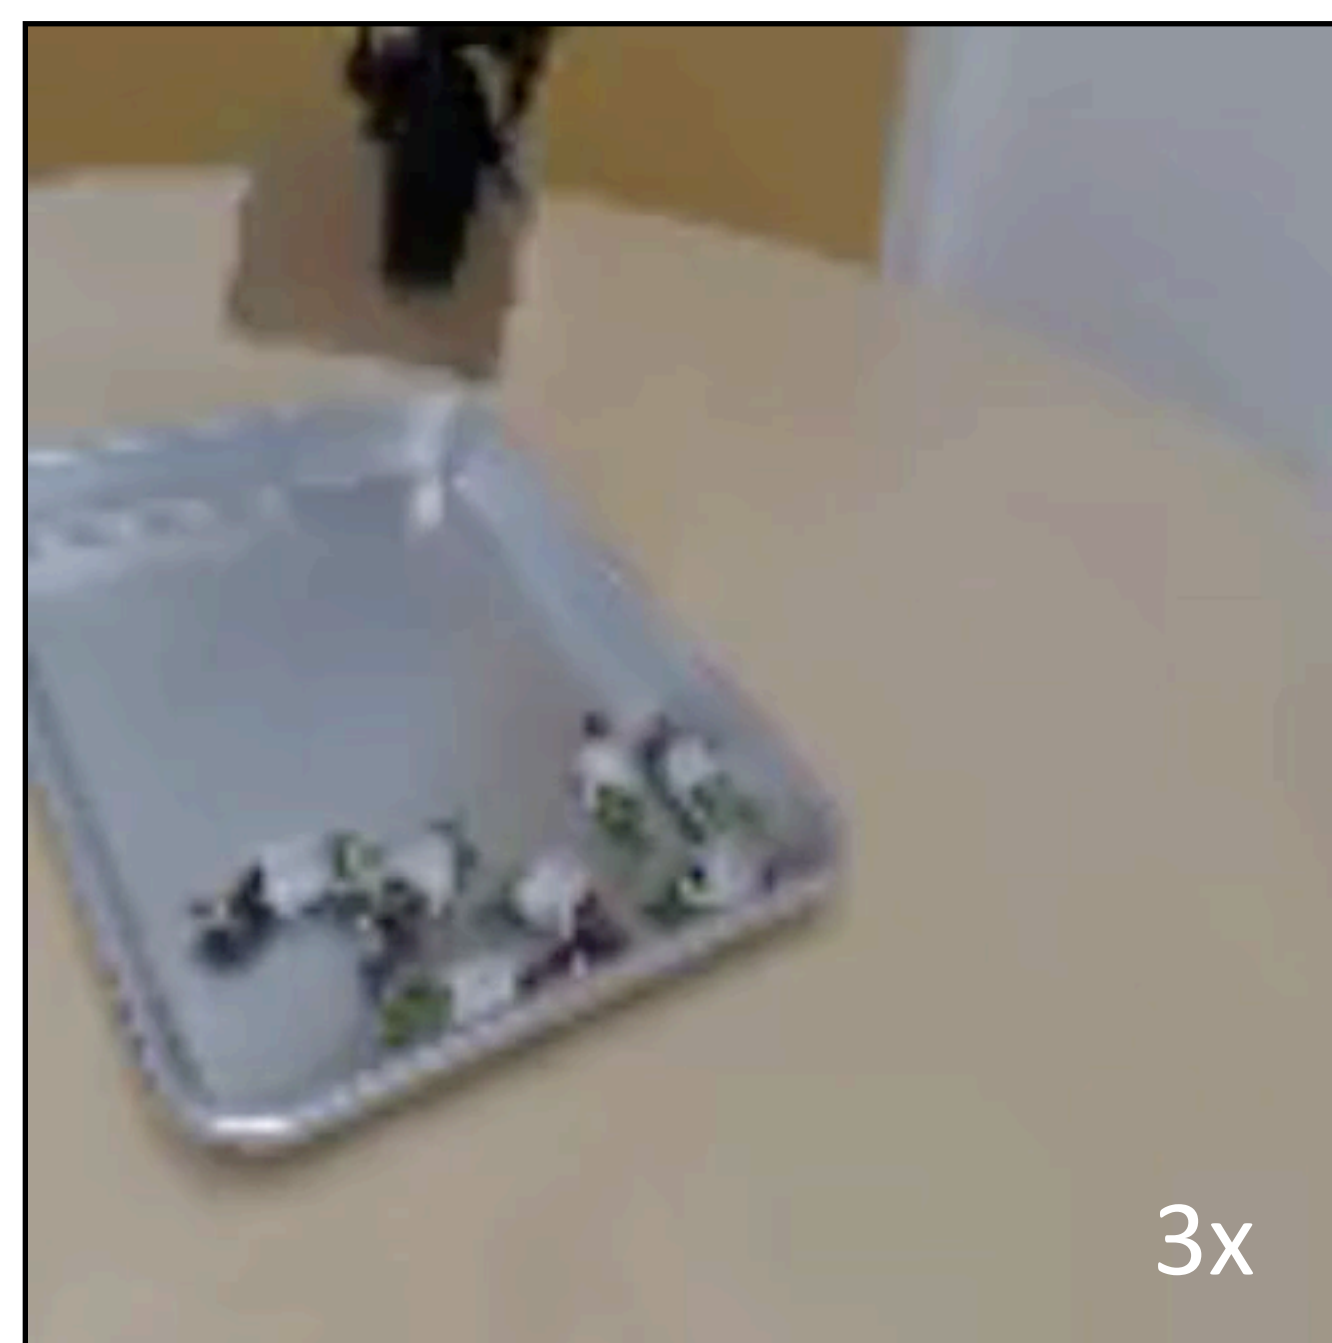


Evaluation (Candy Sweeping)




Behavior Cloning

31%



Ours (without Fine-tuning)

47%



Ours (after 3 Hours of Fine-tuning)

66%

Results on other tasks

Offline BC-> RoboFUME after 3 hrs of fine-tuning



60% -> 85%



45% -> 80%



60% -> 95%




45% -> 55%

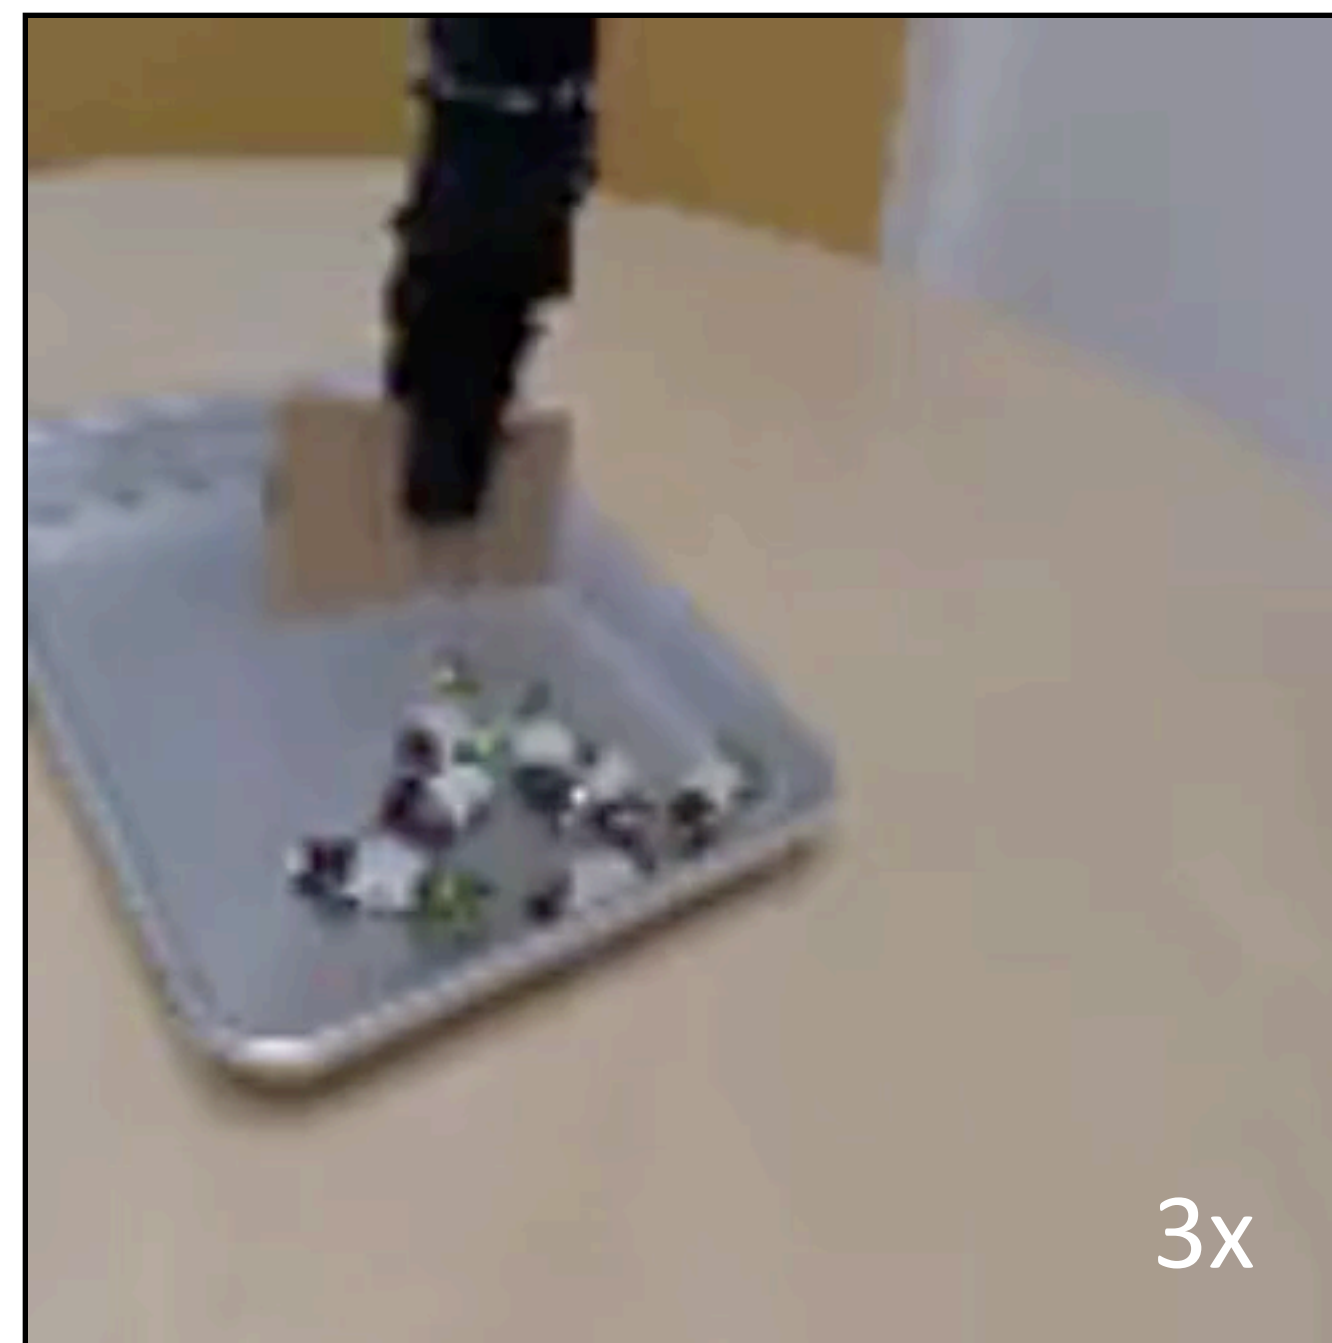
Robot can autonomously collect *useful* data for new tasks.

Analysis: Are Prior Data, Language-Conditioning, & Online Data Important?


Ours (after 3 Hours of Fine-tuning): 66%



Ours (Offline Only)
47%



Ours (Offline Only) w/o Language Conditioning
13%



Ours (Offline Only) w/o Prior Data
23%

More experiments in simulation in the paper!

Data for Robotic Manipulation: Takeaways

1. Large-scale data is starting to become a reality for robotic motor control.
2. We have a path towards scaling data collection, via robot autonomy.

The start of a paradigm shift.

What robot learning research usually does

Collect a new data, use it for a project.

Don't use the data again.

for **each** and **every** project



What ML research currently does


Collect a large dataset **once**
(e.g. ImageNet, Wikitext)

Reuse it many times

Or even better: Reuse a pre-trained model

Let's figure out how to make this the norm!

IRIS Lab



Questions?

