Augmenting Robotic Manipulation through Natural Language

Albert Yu

July 30, 2026

Ideal Home Robot

The Jetsons cartoon (1962-3, 1985-7)

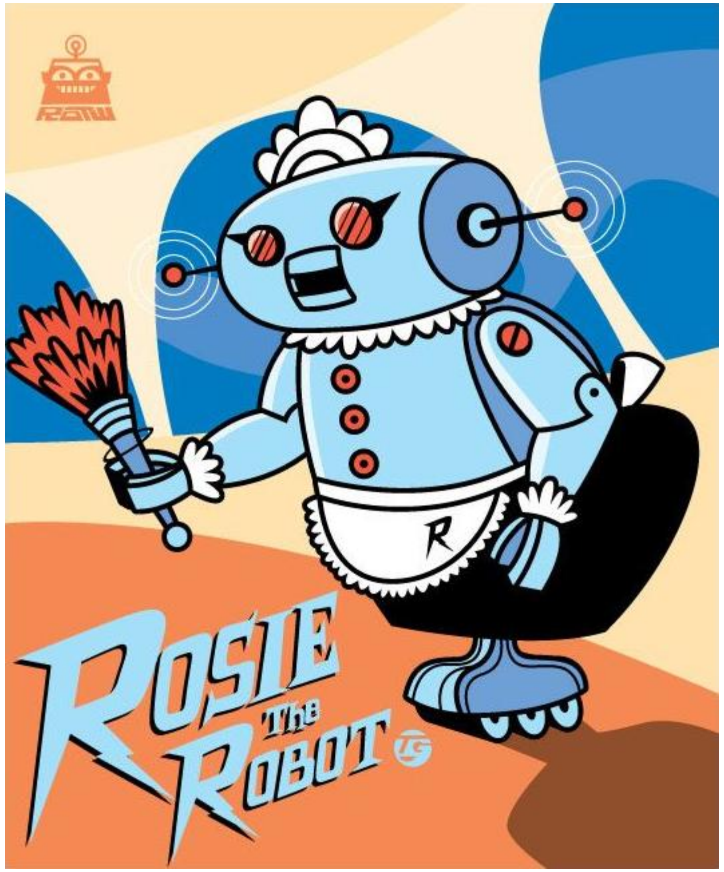

Desired Traits

- ✔ Automates repetitive household tasks
- ✔ Learns new tasks quickly
- ✔ Helps us on complex tasks
- ✔ Follows our instructions + preferences

Desired Traits

- ✔ Automates repetitive household tasks
- ✔ Learns new tasks quickly

Challenge: Needs tons of data

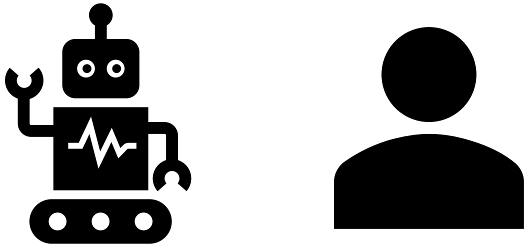

$\sim 1000s$ of demos/task

- ✔ Helps us on complex tasks
- ✔ Follows our instructions + preferences

Challenge: Understanding how to follow/help humans.

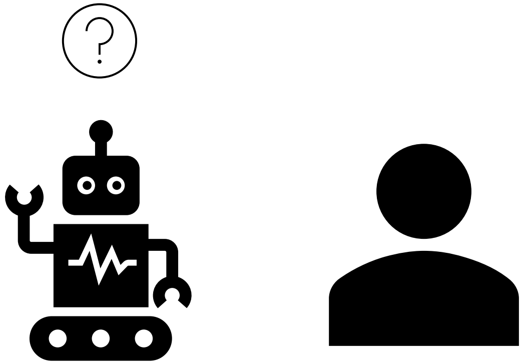

Why might natural language help?

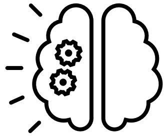

Compact, generalizable knowledge representation

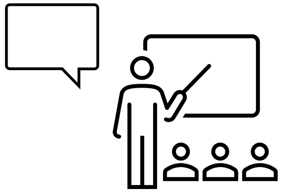

Helps us teach others

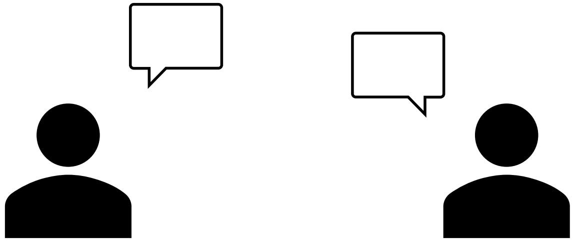

Crucial form of communication

Enables collaboration

Challenge: Needs tons of data

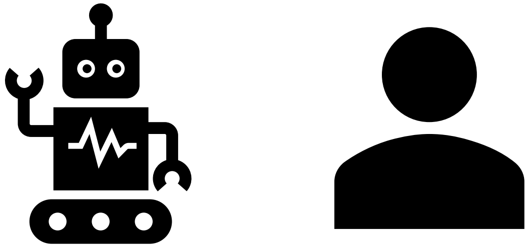

$\sim 1000s$ of demos

Challenge: Understanding how to follow/help humans.

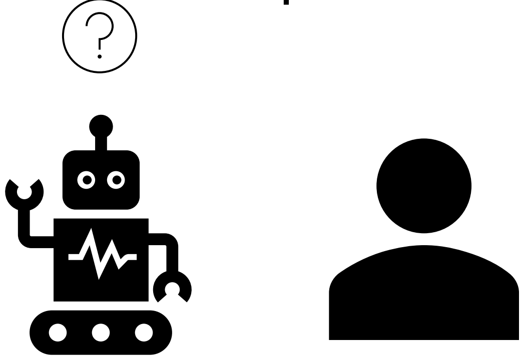

How Language Helps:

Task Instructions for Sample Efficiency

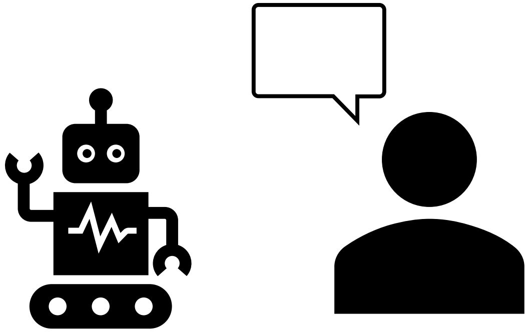

$\sim 10s$ of demos

Communication + Collaboration

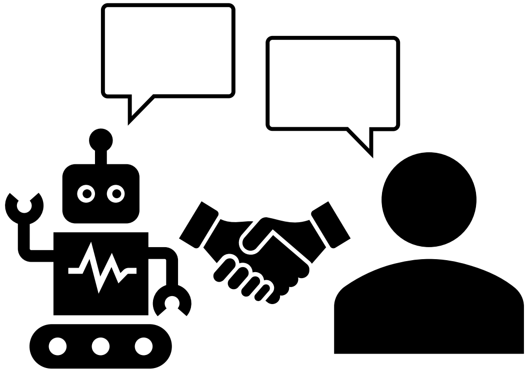

Thesis: Language Improves Robotic Manipulation Capabilities.

Semantics for Generalization

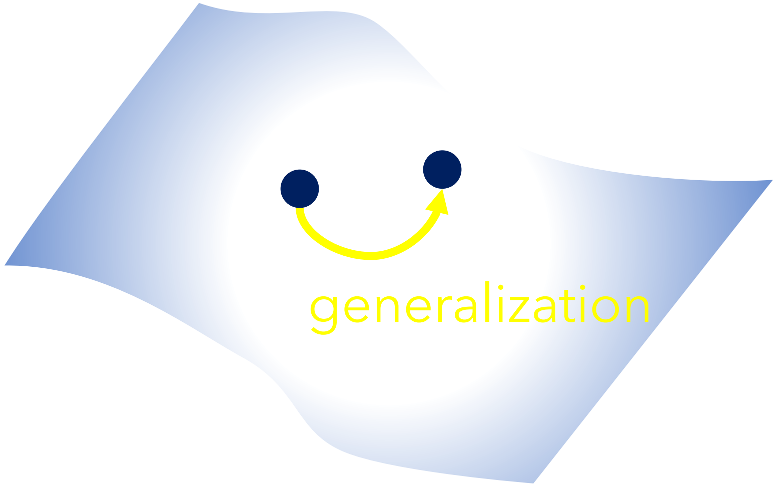

Communicative Medium

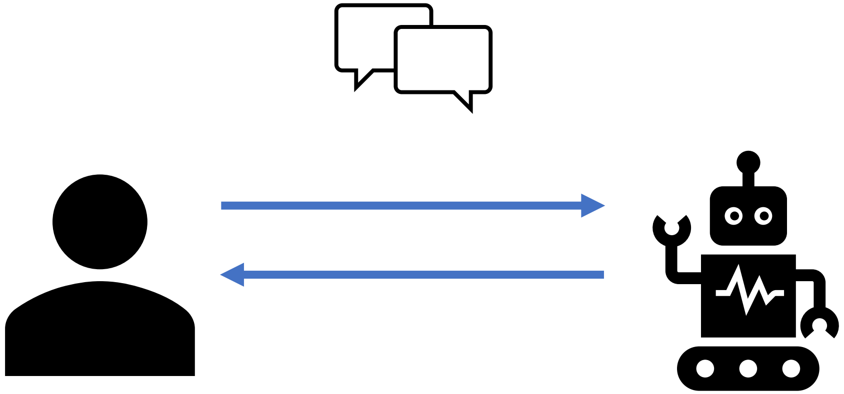

Outline

Language Improves Robotic Capabilities through...

Semantics for Generalization

1. DeL-TaCo: Simultaneous Learning from Demo + Language
2. Lang4Sim2Real: Language to Bridge the Sim2Real Gap

Communicative Medium

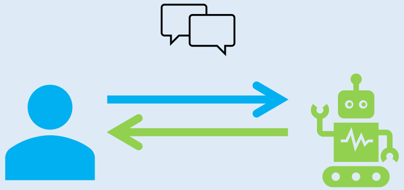

3. MicoBot: Bidirectional Dialog for Human-Robot Collaboration
4. BookWorm: Learning to operate new devices with guidebooks and human feedback

Future Work and Conclusion

Outline

Language Improves Robotic Capabilities through…

Semantics for Generalization

1. **DeL-TaCo**: Simultaneous Learning from Demo + Language
2. Lang4Sim2Real: Language to Bridge the Sim2Real Gap

Communicative Medium

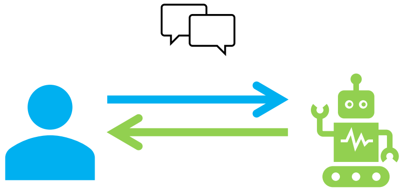

3. MicoBot: Bidirectional Dialog for Human-Robot Collaboration
4. BookWorm: Learning to operate new devices with guidebooks and human feedback

Future Work and Conclusion

Using Both Demonstrations and Language Instructions to Efficiently Learn Robotic Tasks

Collaborators: Ray Mooney

ICLR 2023

Learning new skills

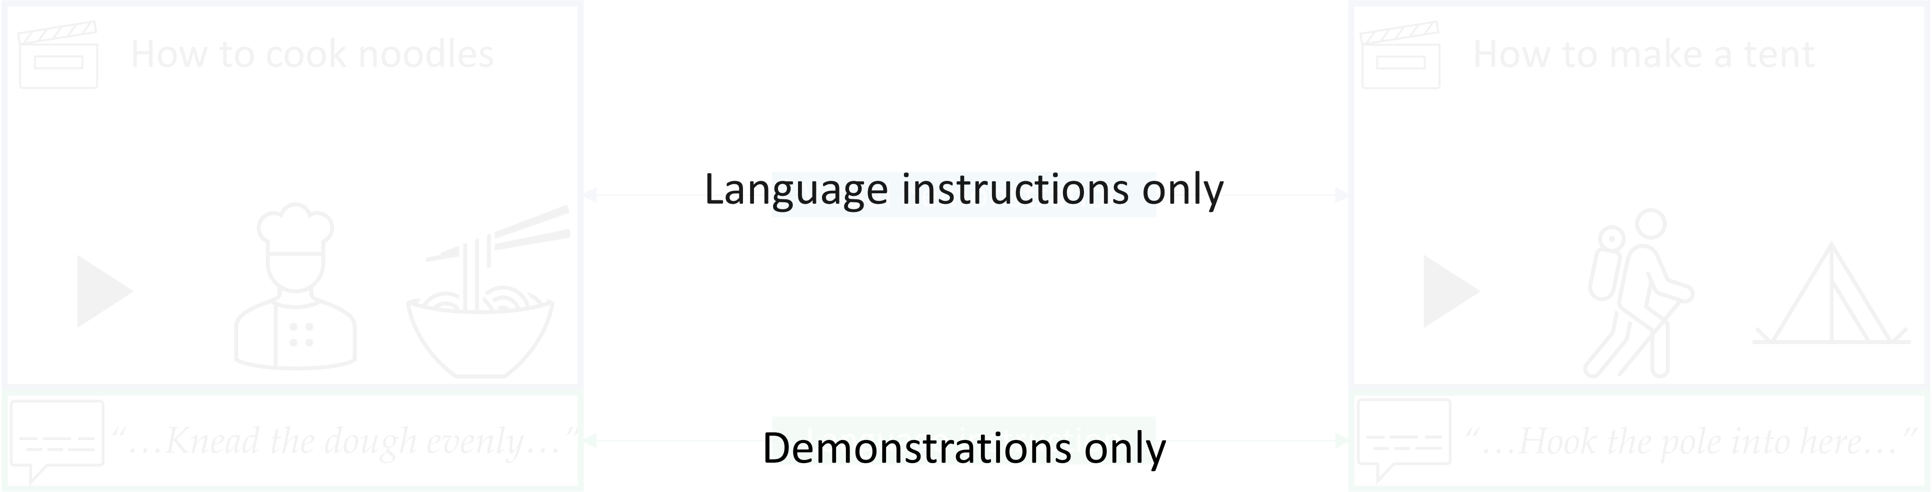

Imagine learning from only one modality…

Hypothesis: Multimodality helps Robots learn new tasks

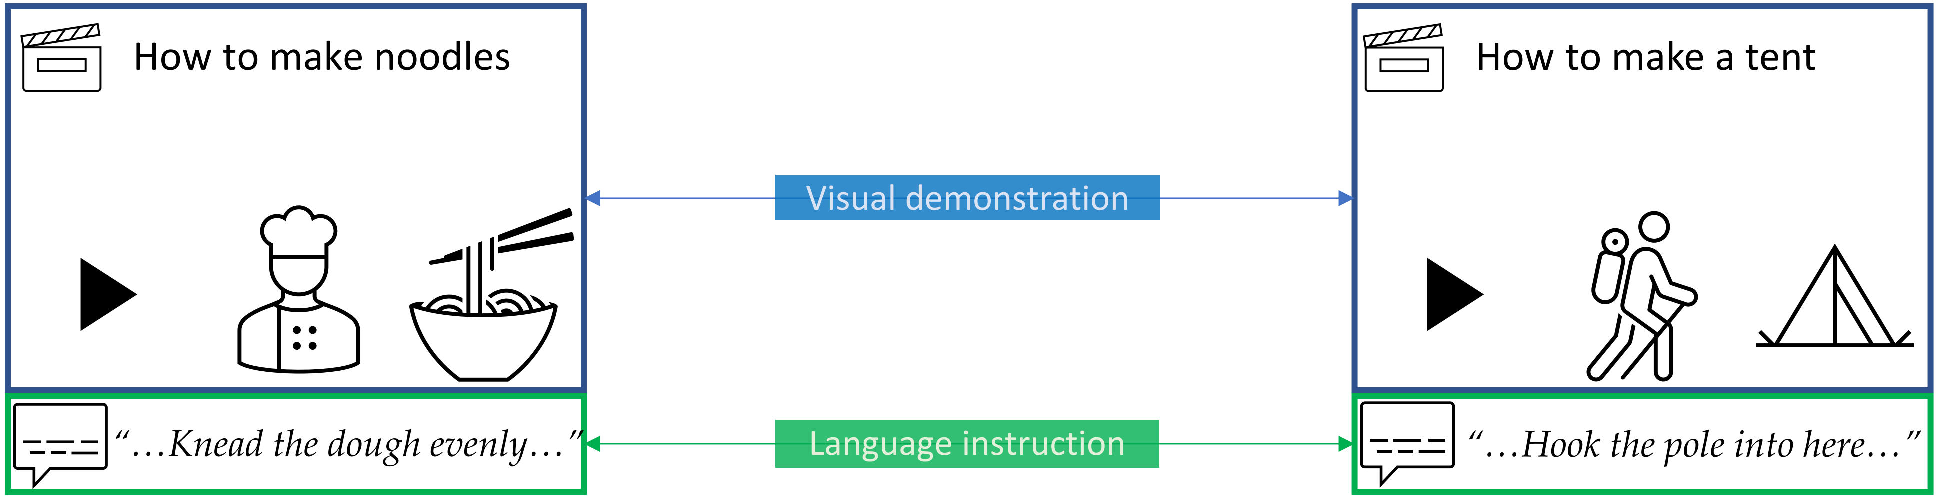

Prior Work

MCIL (RSS 2021)
Lynch & Sermanet

BC-Z (CoRL 2022)
Jang, et. al.

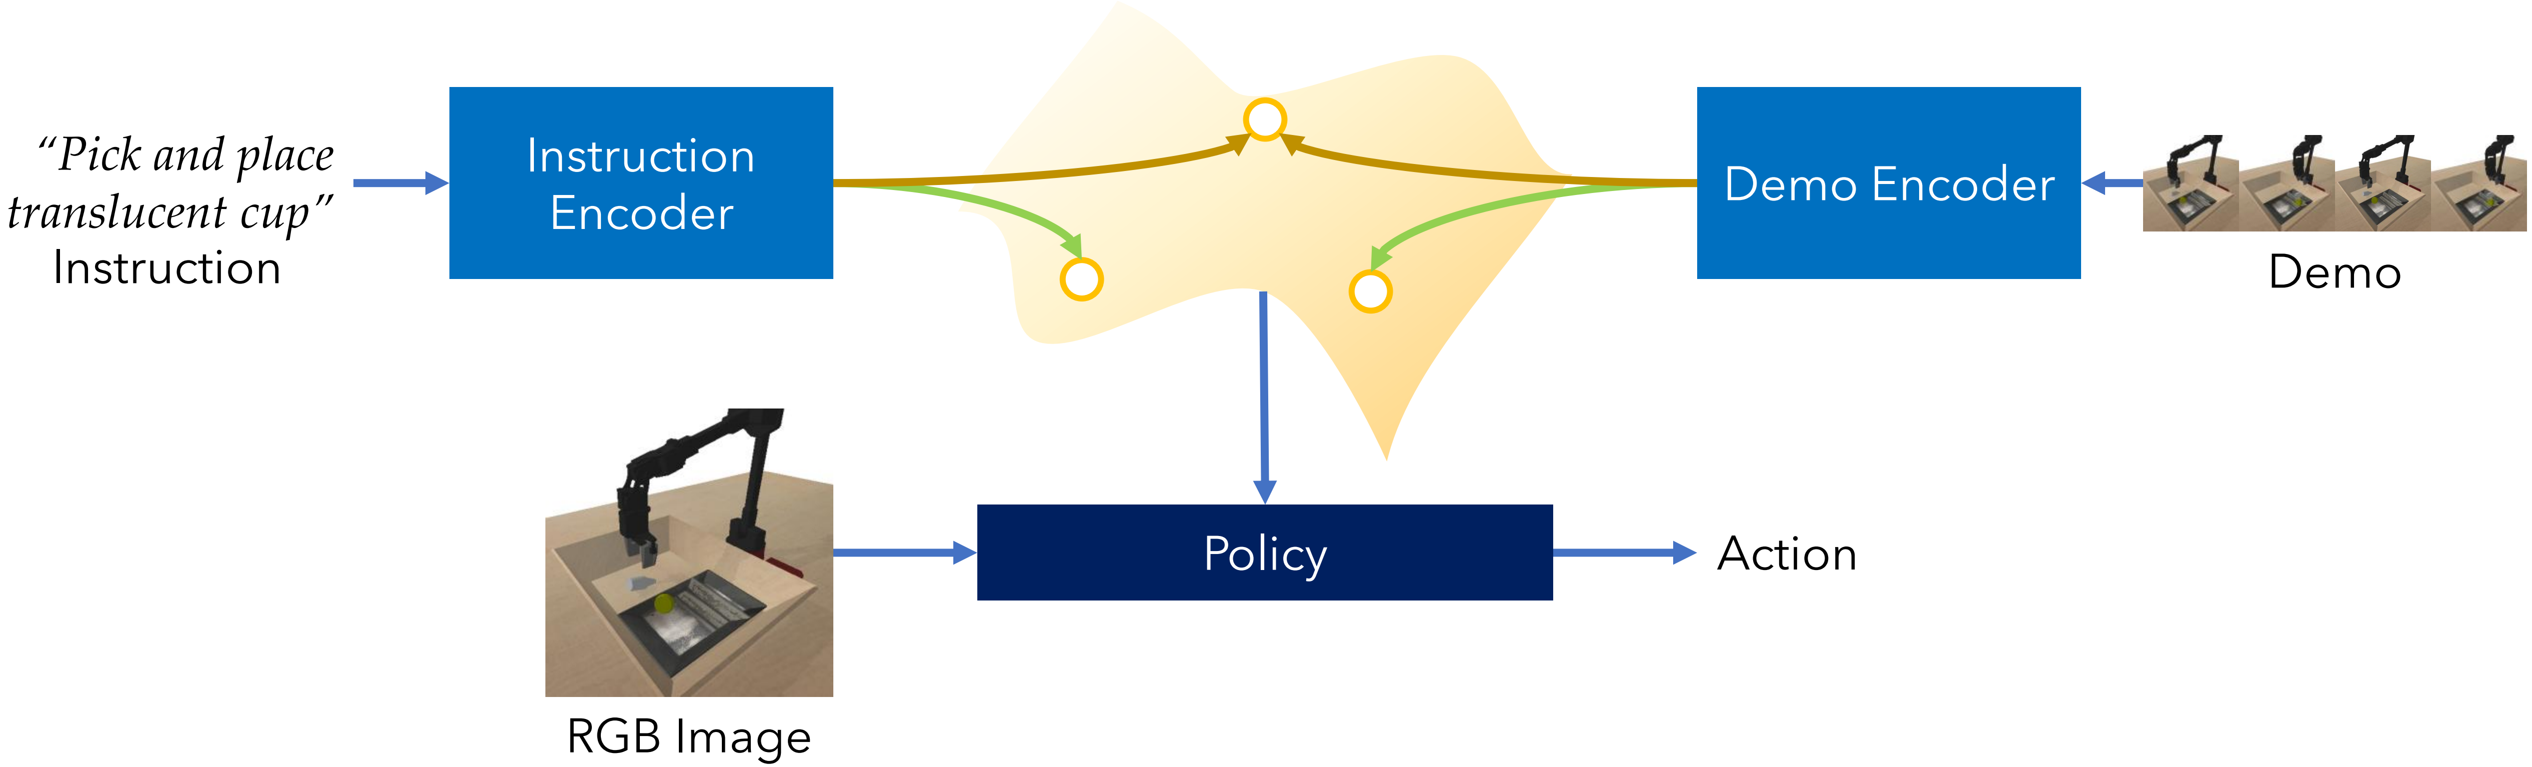

Existing methods do not simultaneously use both demonstrations and language to teach robots new tasks.

MCIL: Corey Lynch, Pierre Sermanet. "Language Conditioned Imitation Learning over Unstructured Data."
BC-Z : Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, Chelsea Finn. "Zero-shot Task Generalization with Robotic Imitation Learning."

DeL-TaCo

Joint Demo-Language Task Conditioning

Leverage benefits of learning with **both demonstrations and language**

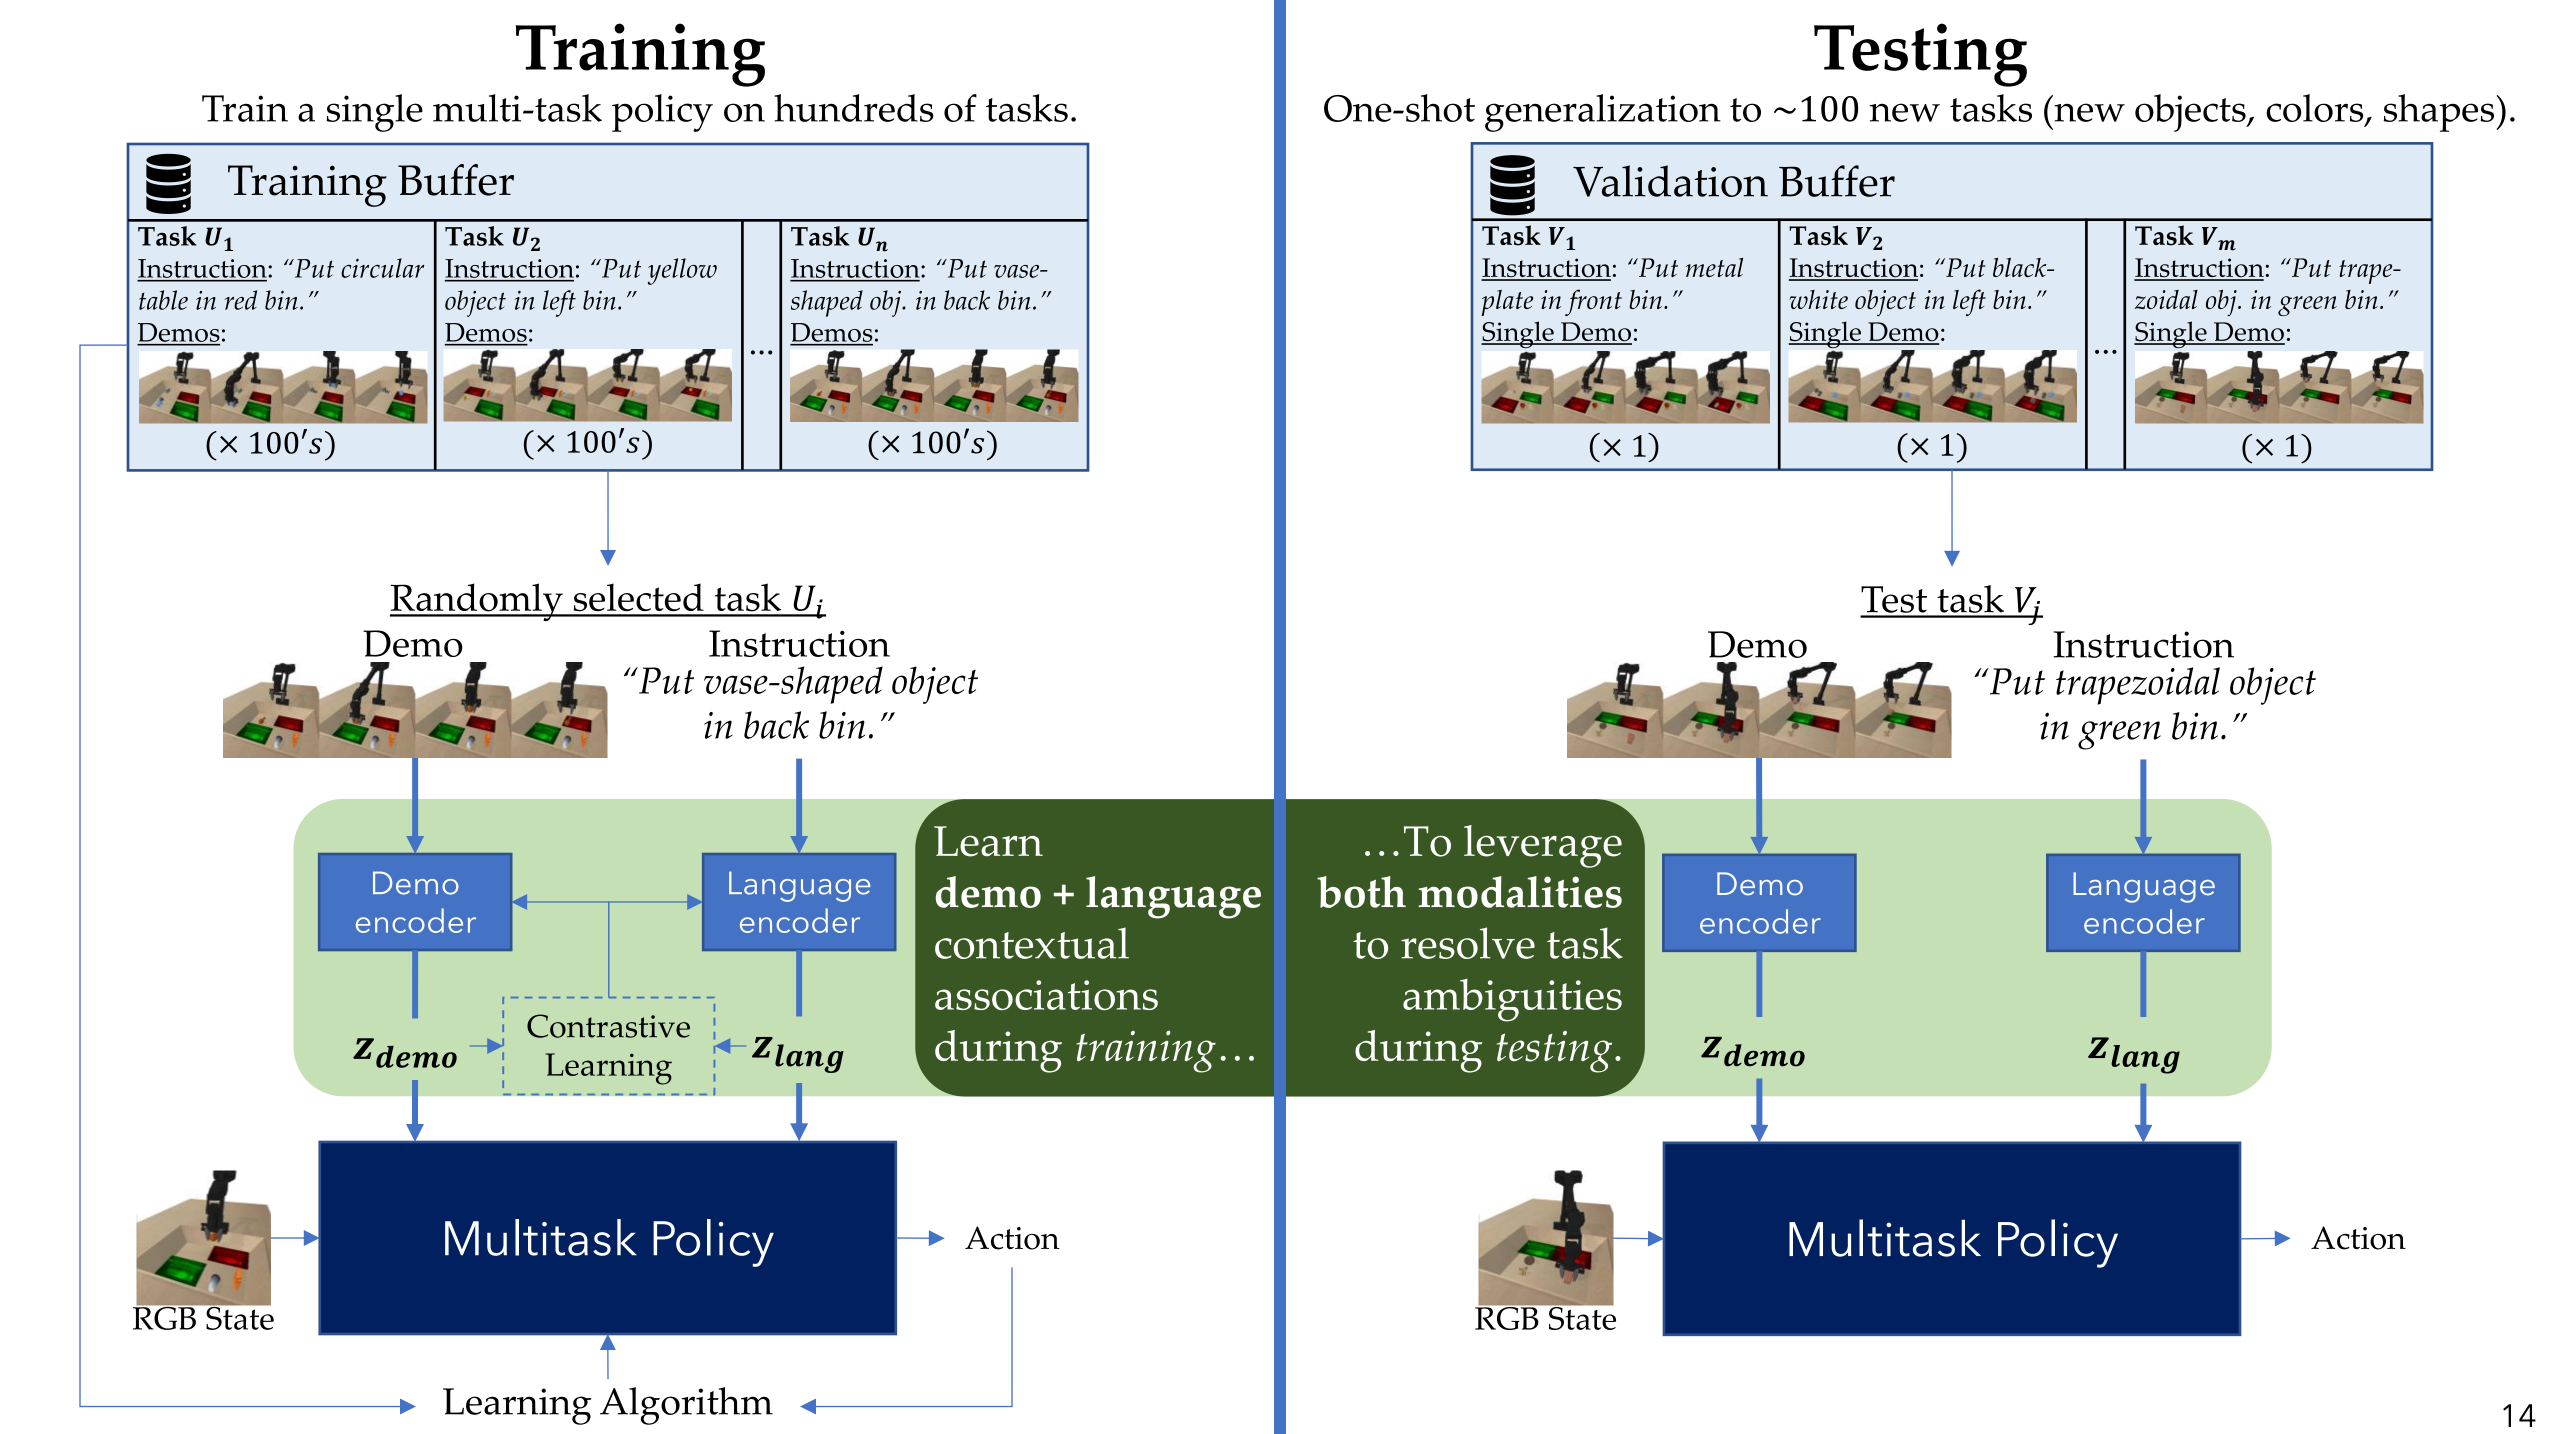
Training
Train a single multi-task policy on hundreds of tasks.
Training Buffer
Task U_1
Instruction: "Put circular table in red bin."
Demos:
(× 100's)
Task U_2
Instruction: "Put yellow object in left bin."
Demos:
(× 100's)
…
Task U_n
Instruction: "Put vase-shaped obj. in back bin."
Demos:
(× 100's)
Randomly selected task U_i
Demo
Instruction
"Put vase-shaped object in back bin."
Demo encoder
Language encoder
z_{demo}
Contrastive Learning
z_{lang}
Learn **demo + language** contextual associations during *training*…
Multitask Policy
RGB State
Action
Learning Algorithm
Testing
One-shot generalization to ~100 new tasks (new objects, colors, shapes).
Validation Buffer
Task V_1
Instruction: "Put metal plate in front bin."
Single Demo:
(× 1)
Task V_2
Instruction: "Put black-white object in left bin."
Single Demo:
(× 1)
…
Task V_m
Instruction: "Put trape-zoidal obj. in green bin."
Single Demo:
(× 1)
Test task V_j
Demo
Instruction
"Put trapezoidal object in green bin."
…To leverage **both modalities** to resolve task ambiguities during *testing*.
Demo encoder
Language encoder
z_{demo}
z_{lang}
Multitask Policy
RGB State
Action

Example Train Task
Task ID 1: *"Put fountain vase in green bin."*

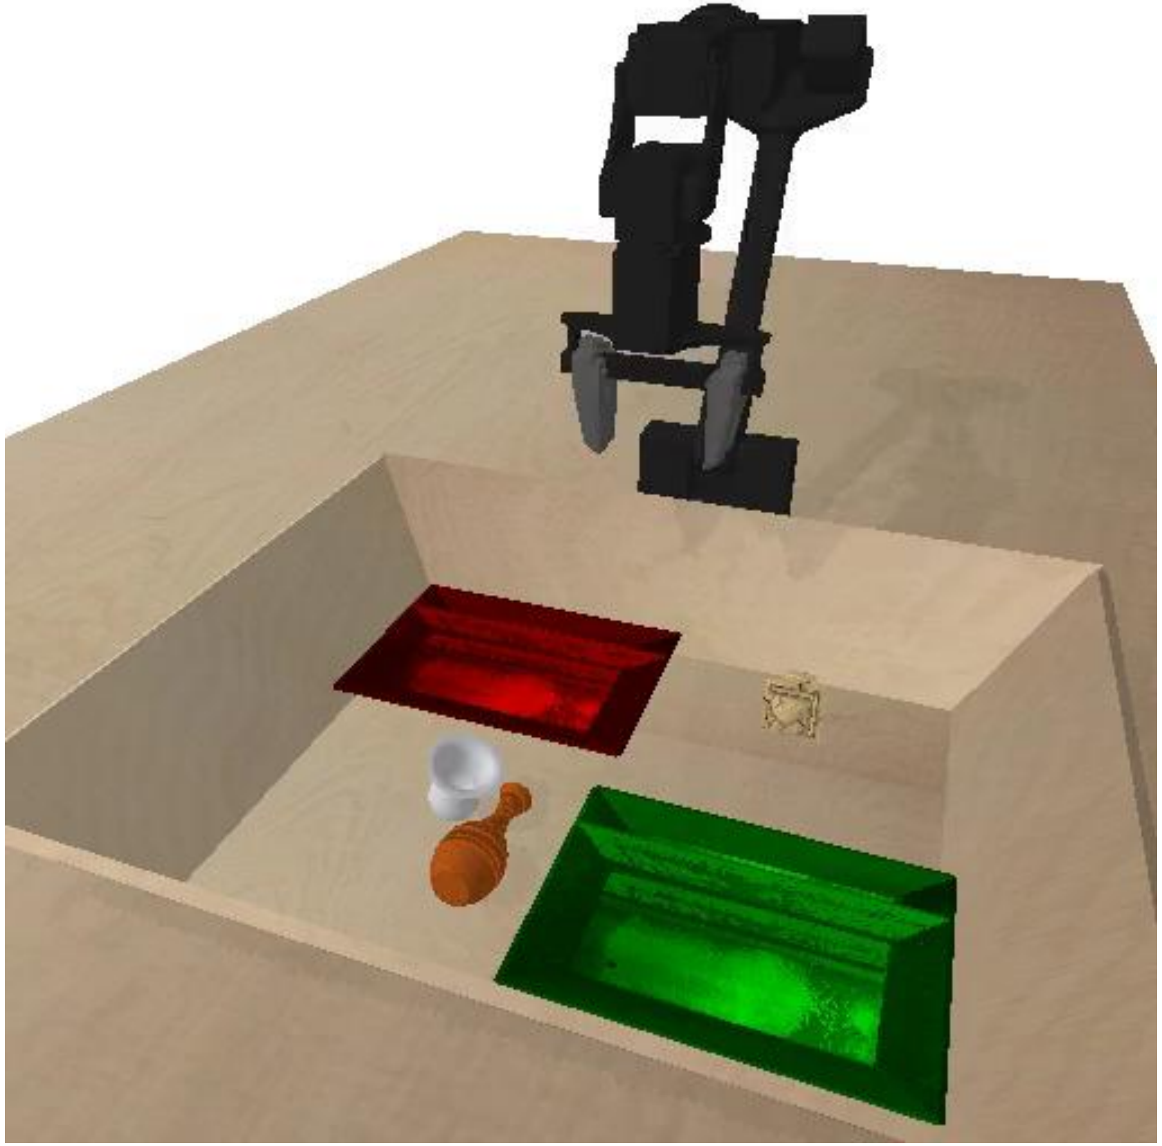

Example Test Task
Task ID 178: *"Put box sofa in back bin."*

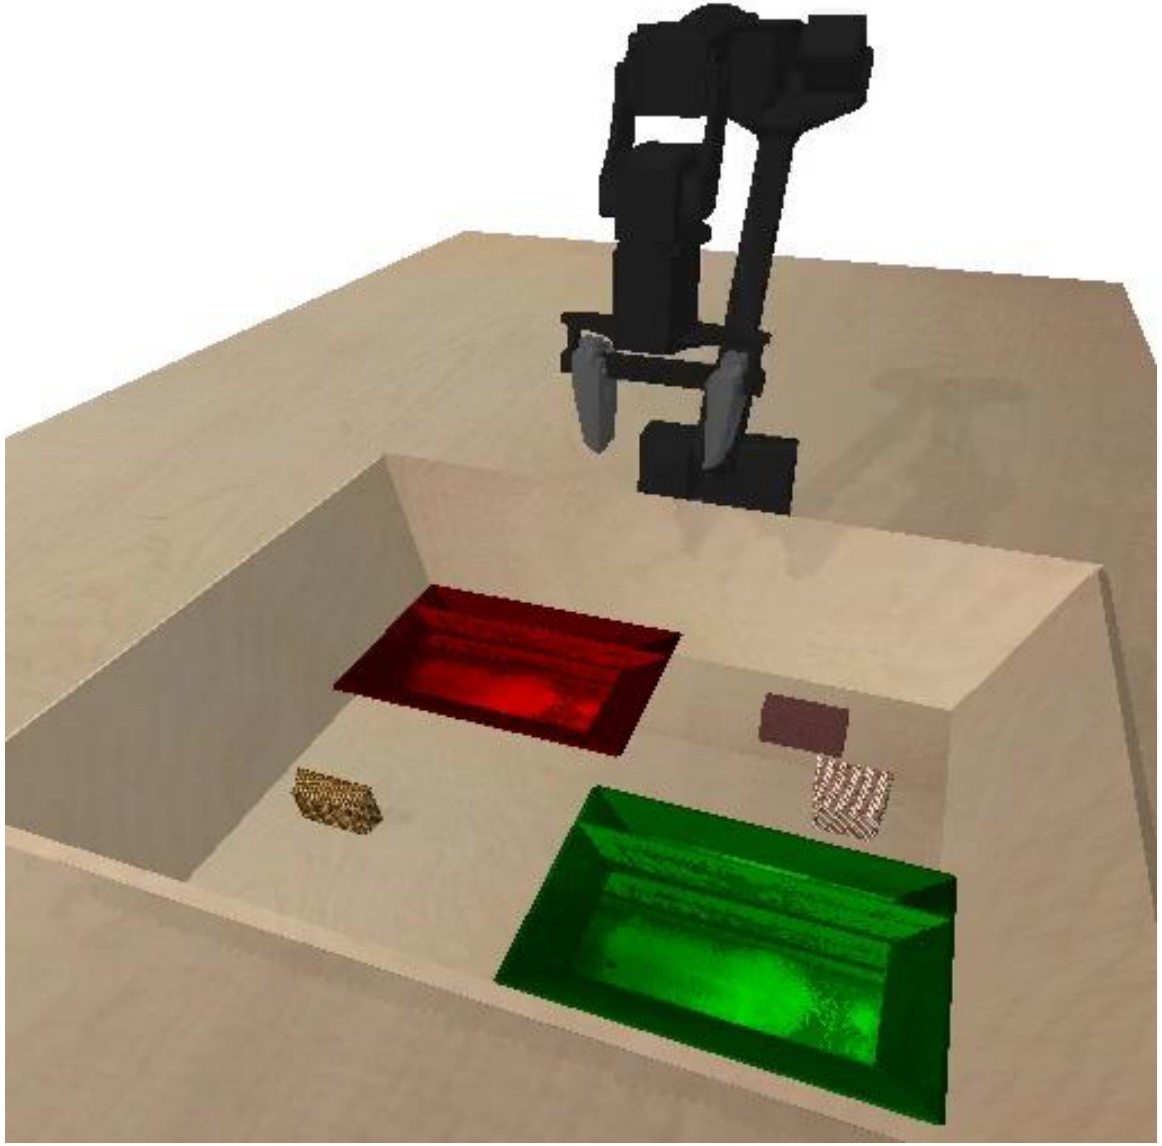

Generalization Performance on Novel Objects, Colors, Shapes

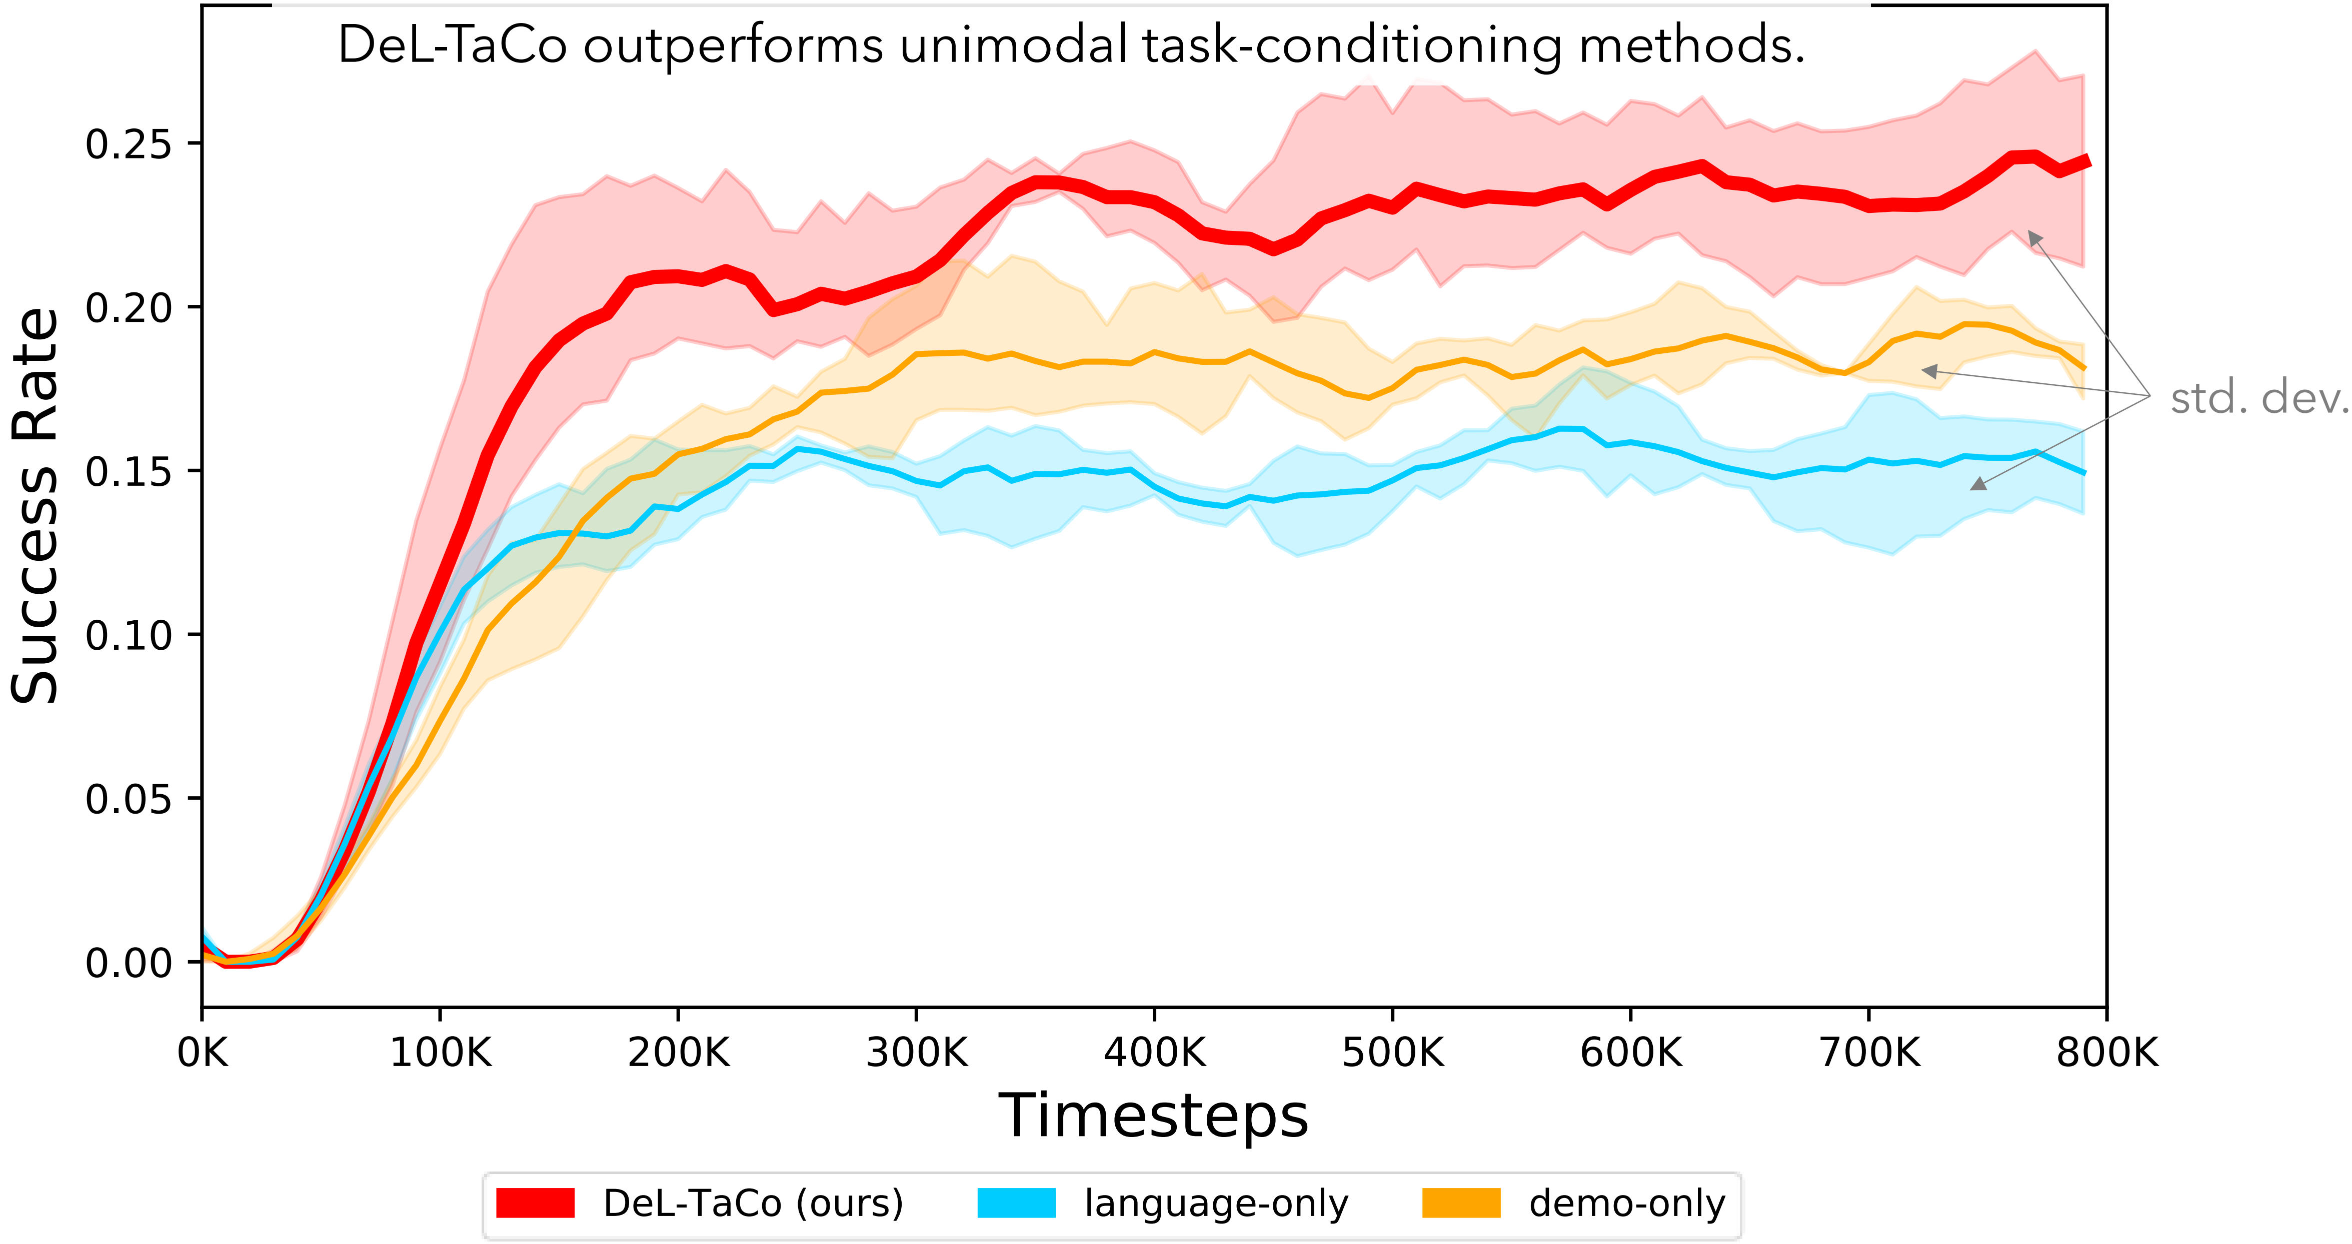

3 seeds per line

How Many Demonstrations is Language Worth?

Providing language with Del-Taco = Finetuning on **25-50 extra demos** per test task.

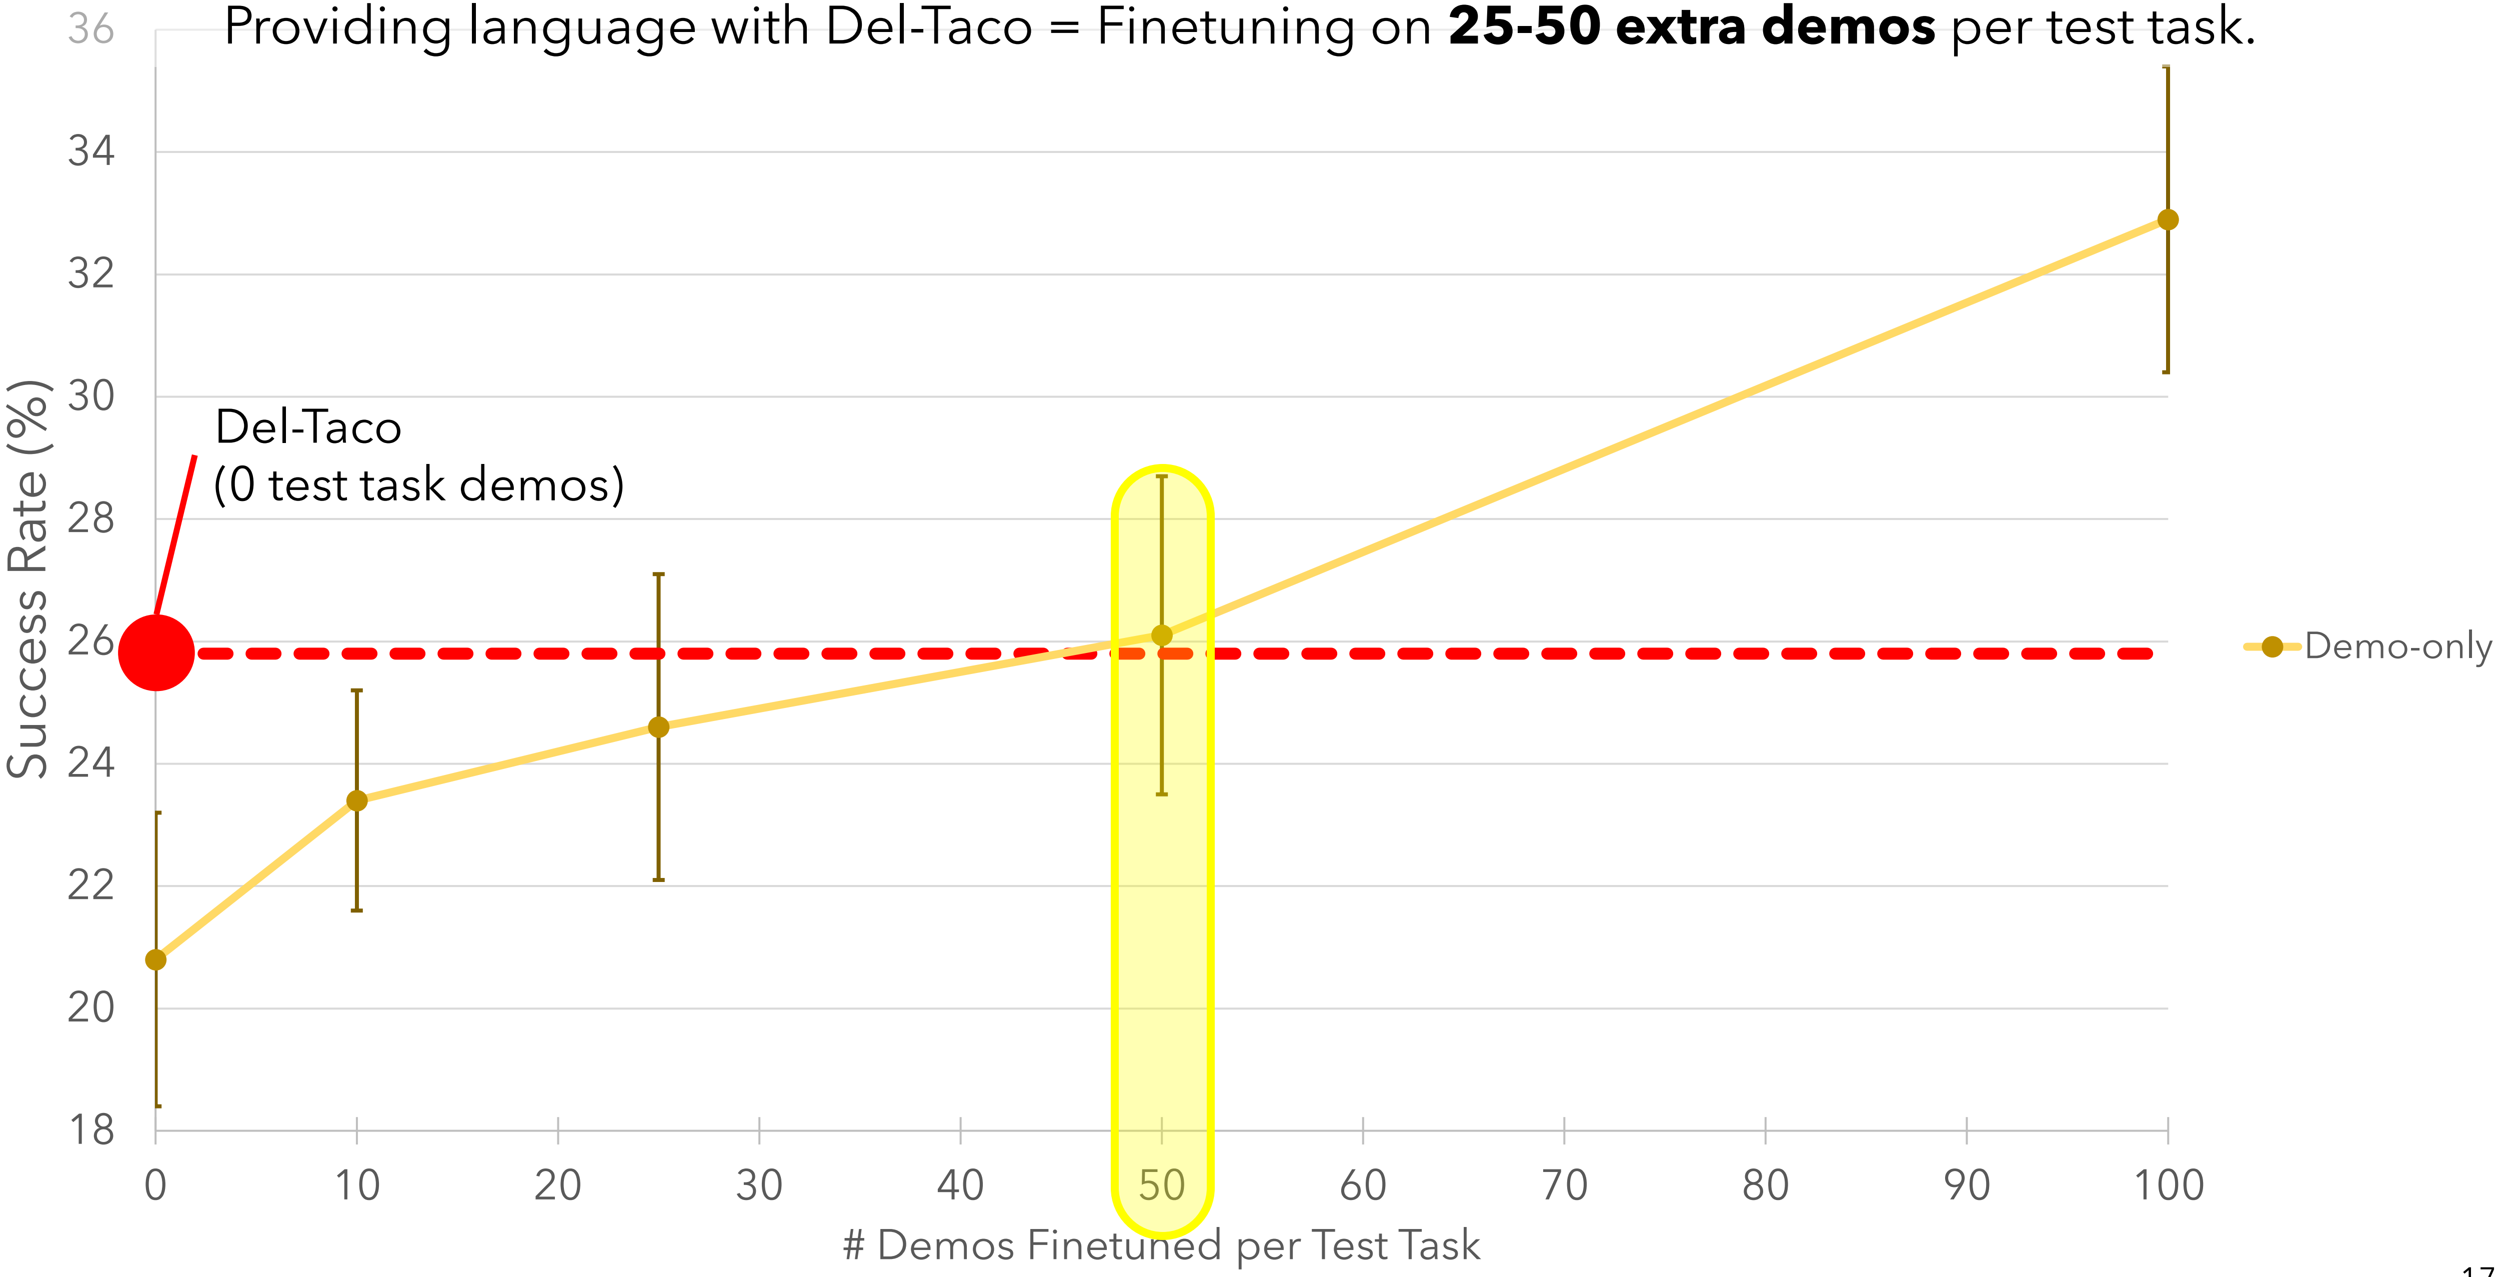

Key Takeaways

Generalization: Learning from both demonstrations and language with DeL-TaCo leads to better generalization than learning from one modality.

Sample-Efficiency: DeL-TaCo reduces the need for ~50 demonstrations over policies conditioned only on demonstrations.

Outline

Language Improves Robotic Capabilities through...

Semantics for Generalization

1. DeL-TaCo: Simultaneous Learning from Demo + Language
2. **Lang4Sim2Real**: Language to Bridge the Sim2Real Gap

Communicative Medium

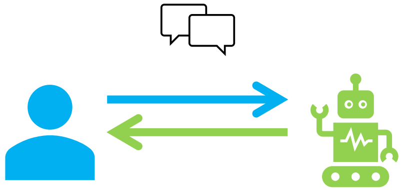

3. MicoBot: Bidirectional Dialog for Human-Robot Collaboration
4. BookWorm: Learning to operate new devices with guidebooks and human feedback

Future Work and Conclusion

Natural Language Can Help Bridge the Sim2Real Gap

Collaborators: Ray Mooney, Roberto Martín-Martín, Addie Foote

RSS 2024

Hypothesis: Semantic information in language enables robots to bridge wide visual sim2real gaps.

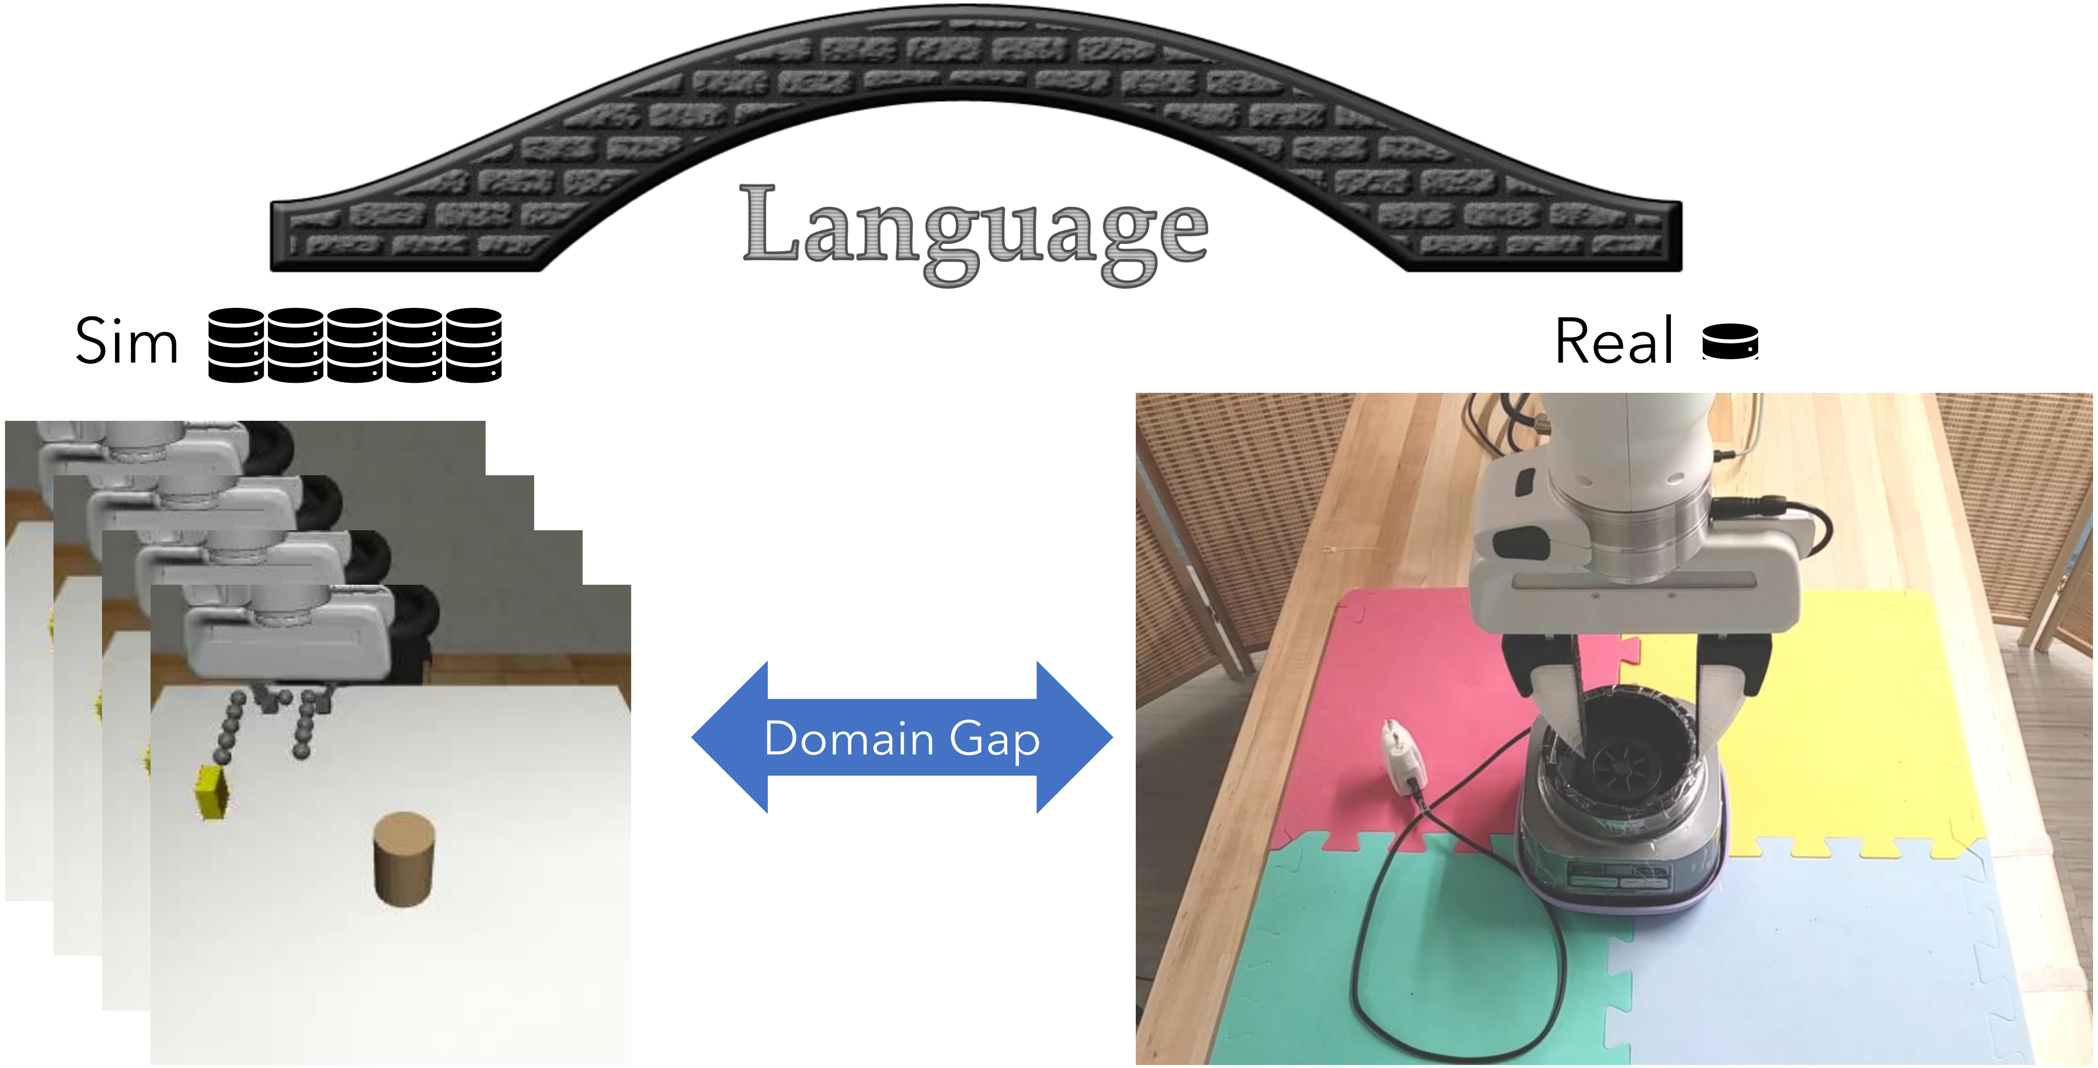

Semantically Similar States → Similar Actions

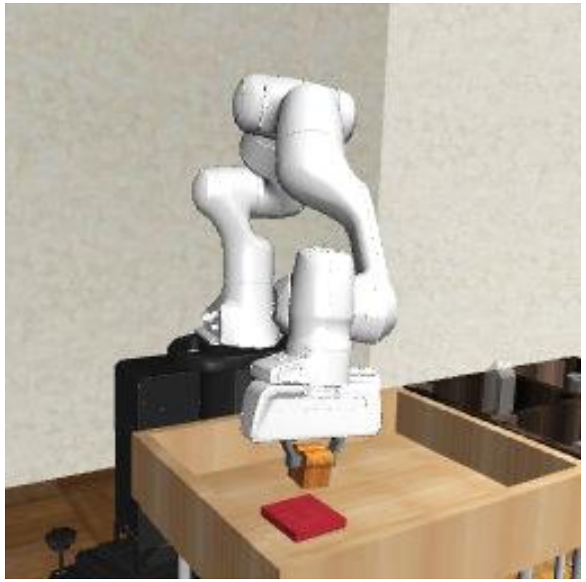

"gripper holding bread above coaster."

"gripper holding carrot above yellow square."

From both states, the robot should open the gripper to place the object.

Domain Invariant Representations

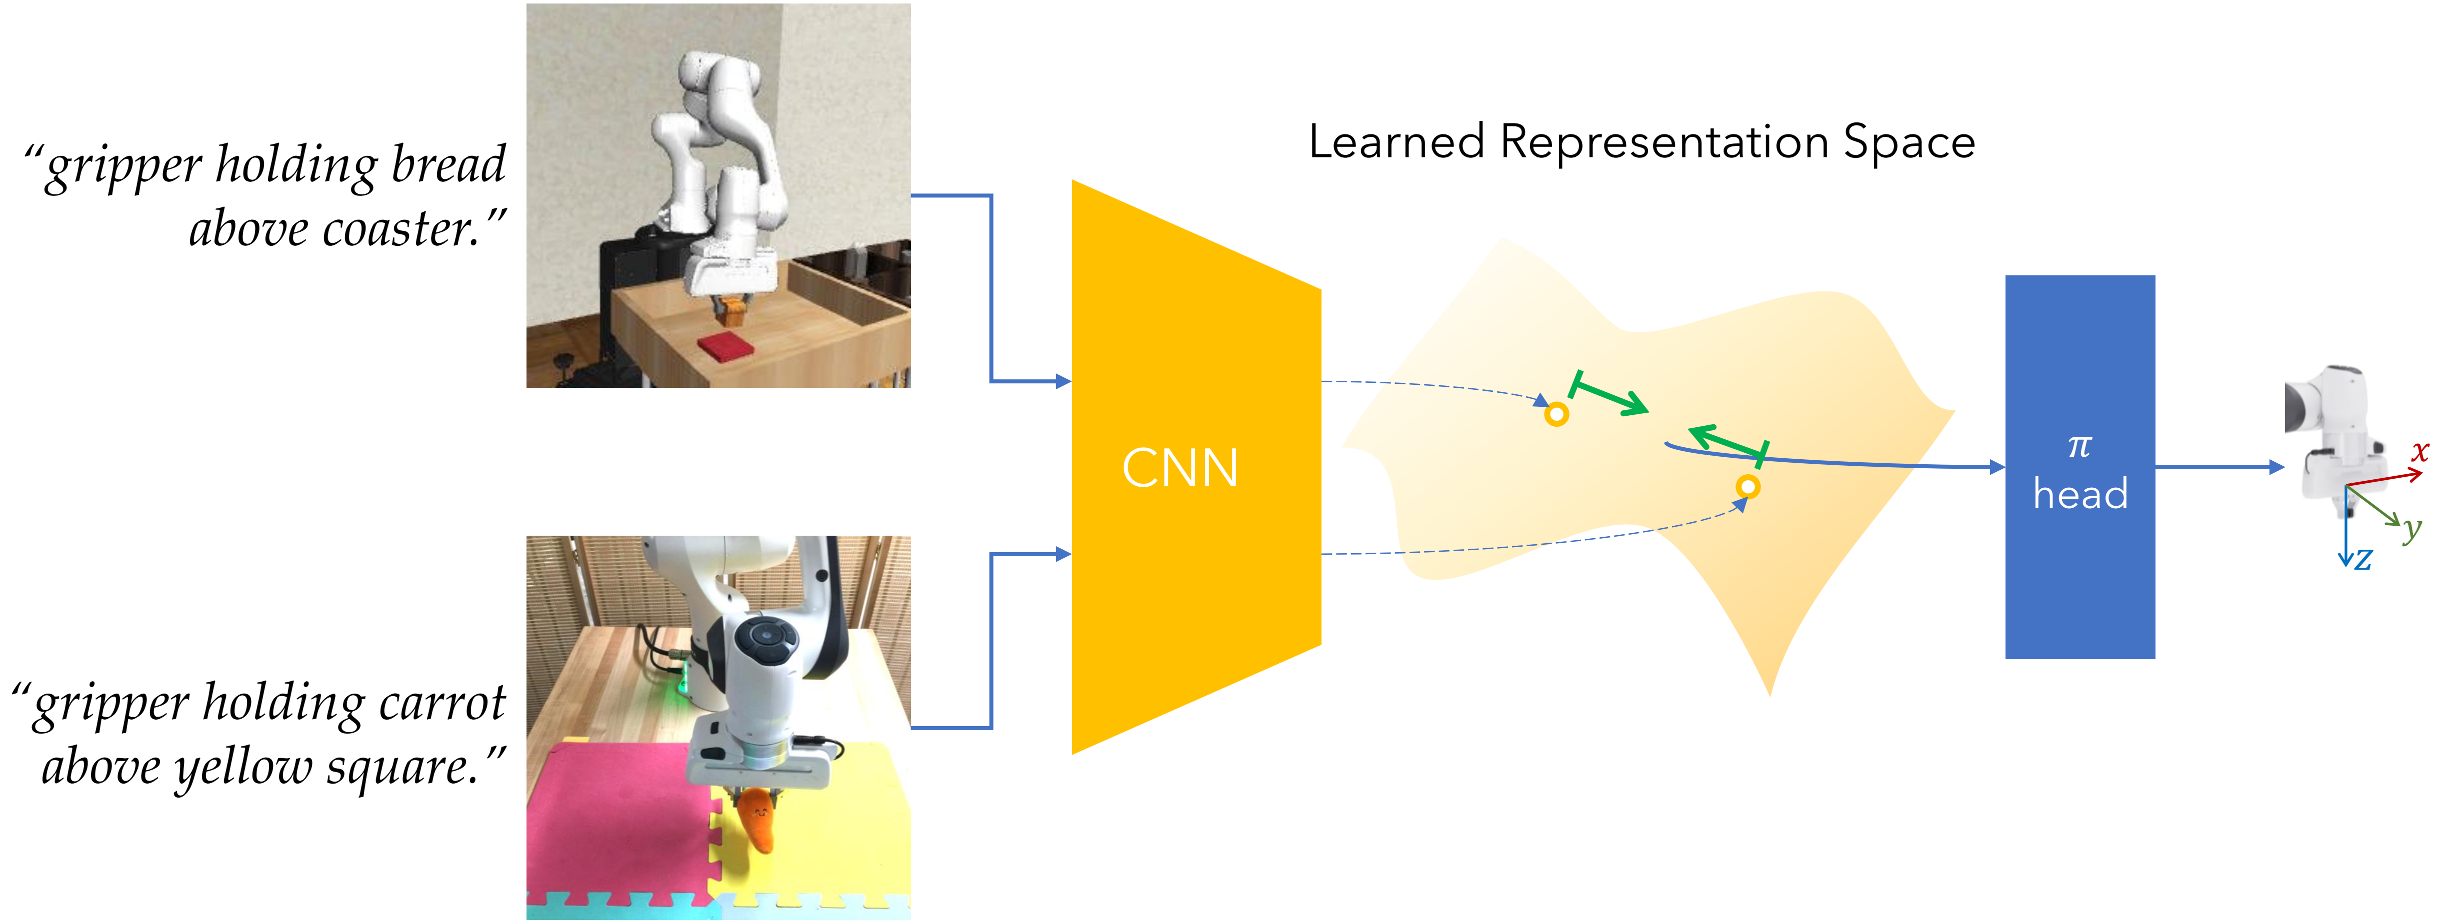

Domain Invariant Representations

"gripper holding milk next to coaster."

"gripper wrapping blender wire."

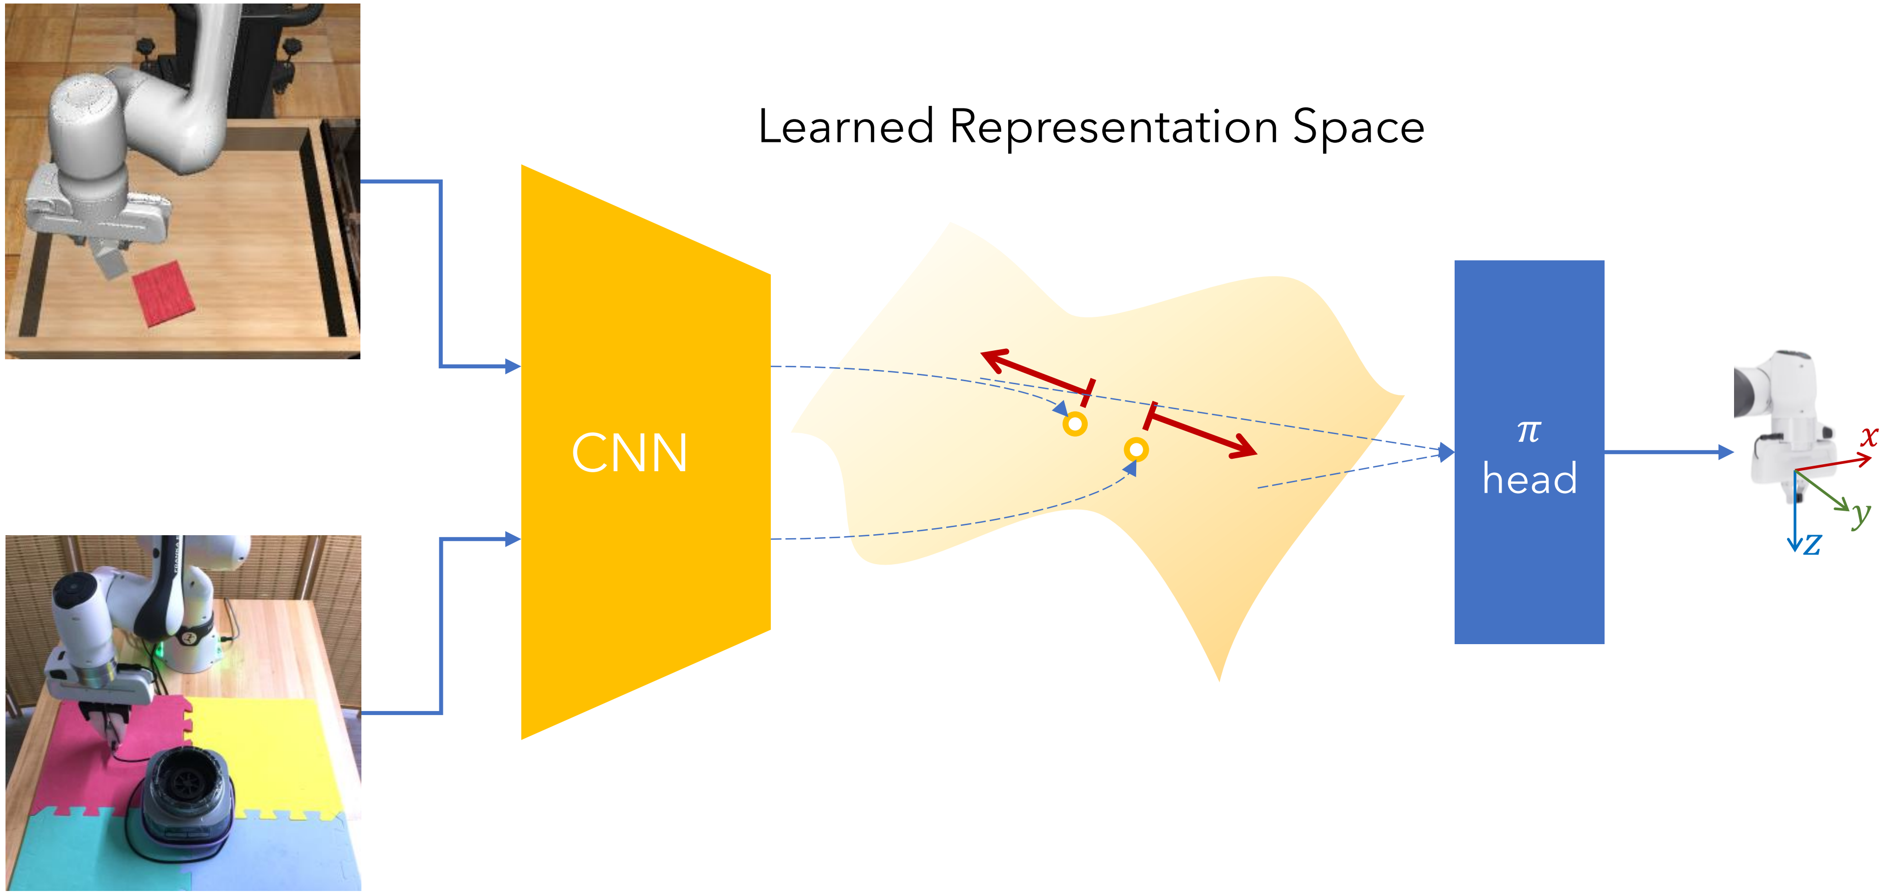

Prior Work: Sim2Real

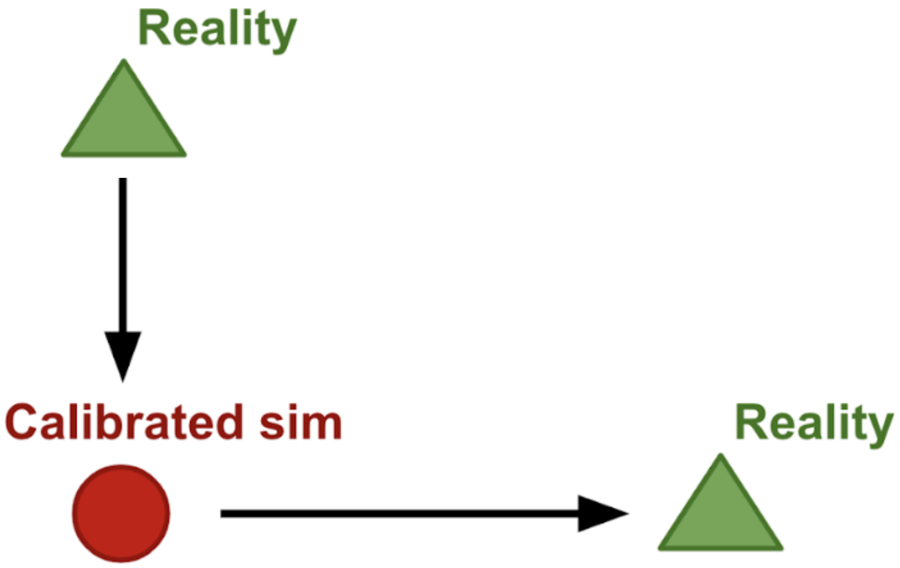

System Identification (Yu+ 2017)

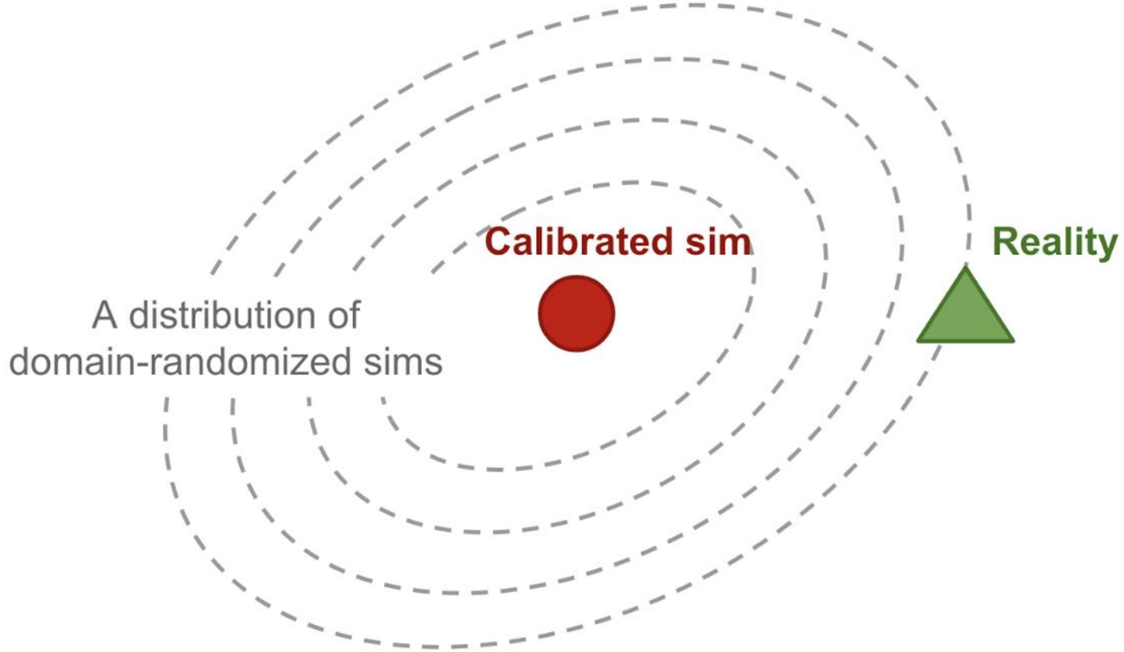

Domain Randomization (Tobin+ 2017)

Both approaches are **very engineering-intensive**.
Domain Randomization leads to overly conservative policies.

https://lilianweng.github.io/posts/2019-05-05-domain-randomization/
Yu+2017: Wenhao Yu, Jie Tan, C Karen Liu, and Greg Turk. "Preparing for the unknown: Learning a universal policy with online system identification." RSS 2017.
Tobin+2017: Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. "Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World." IROS 2017.

Prior Work: Visual Pretraining

Vision-only Pretraining
Image Classification Masked Image Modeling/Reconstruction Contrastive Learning (via img aug) Frame Temporal Ordering Future Frame Prediction (Radosavovic et al., 2023; Zhao et al., 2022; Gupta et al., 2022b; Seo et al., 2023; Laskin et al., 2020; He et al., 2020; Jing et al., 2023; Zhao et al., 2022; Yuan et al., 2022; Wang et al., 2022)

These visual representations are not robust to large domain shifts.
Our work is the first to apply vision-language pretraining for sim2real.

Setup + Approach

① Pretrain Visual Backbone

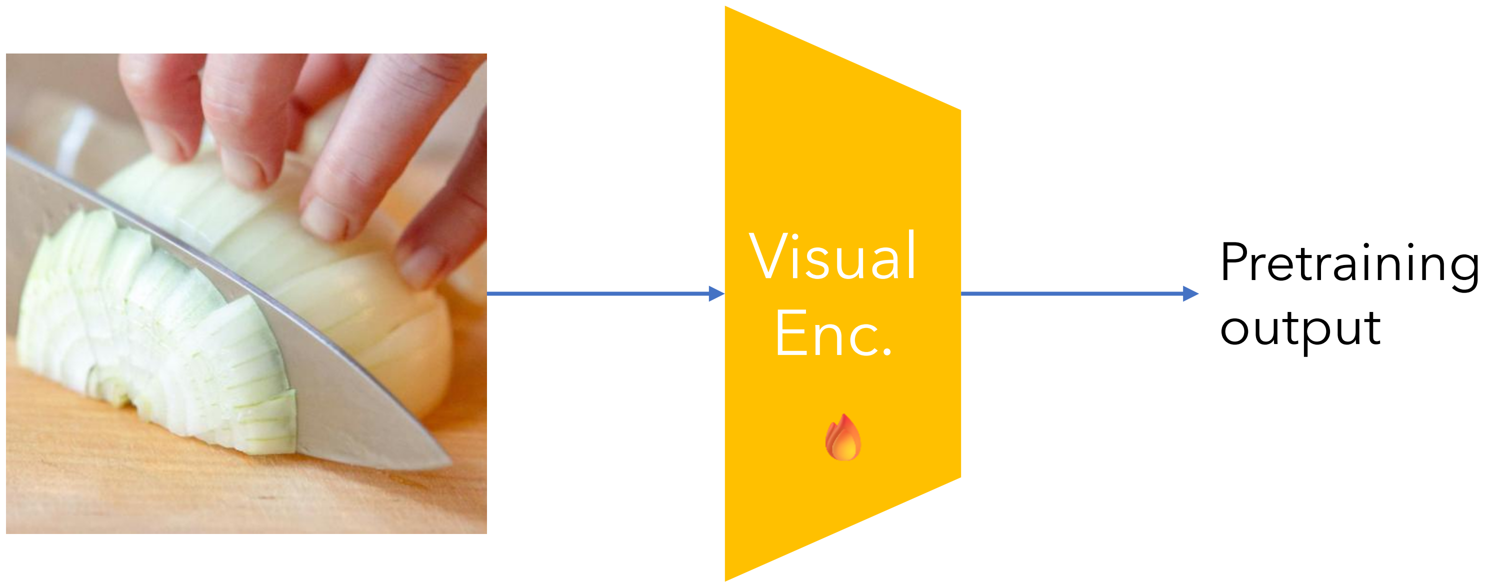

② Finetune Policy Head

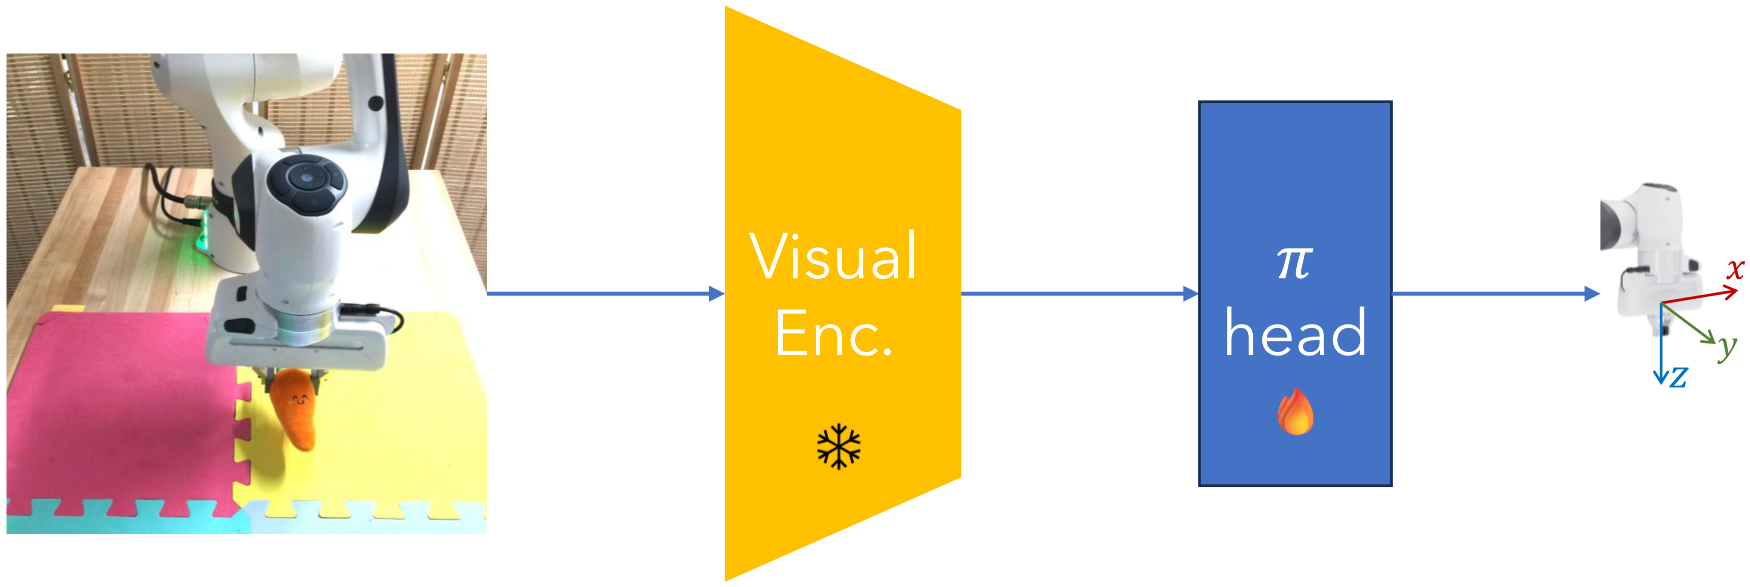

Lang4Sim2Real: Image-Language Pretraining

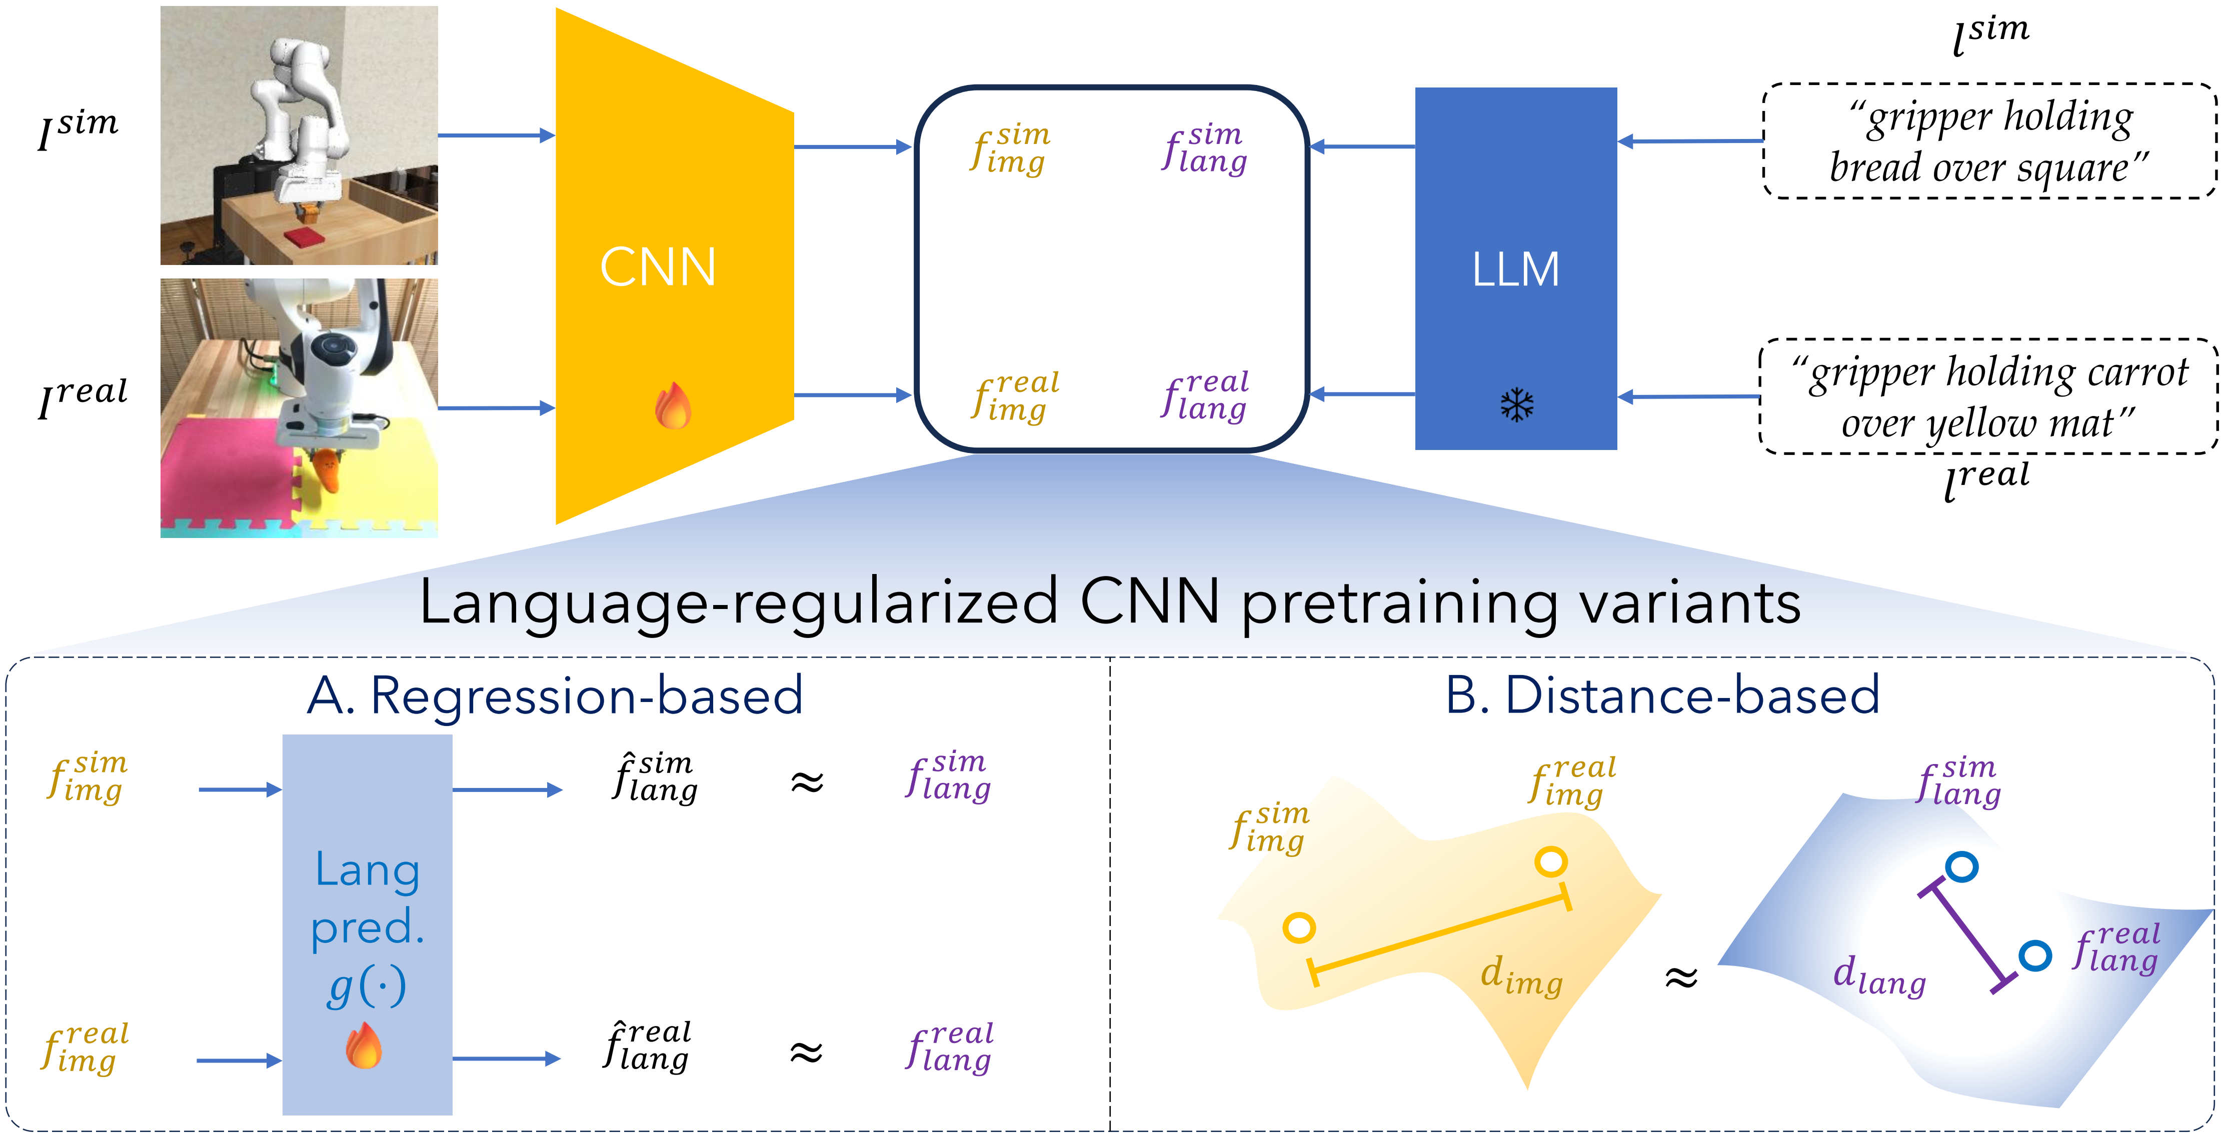

Experiments

Stack Object

2-step Pick and Place

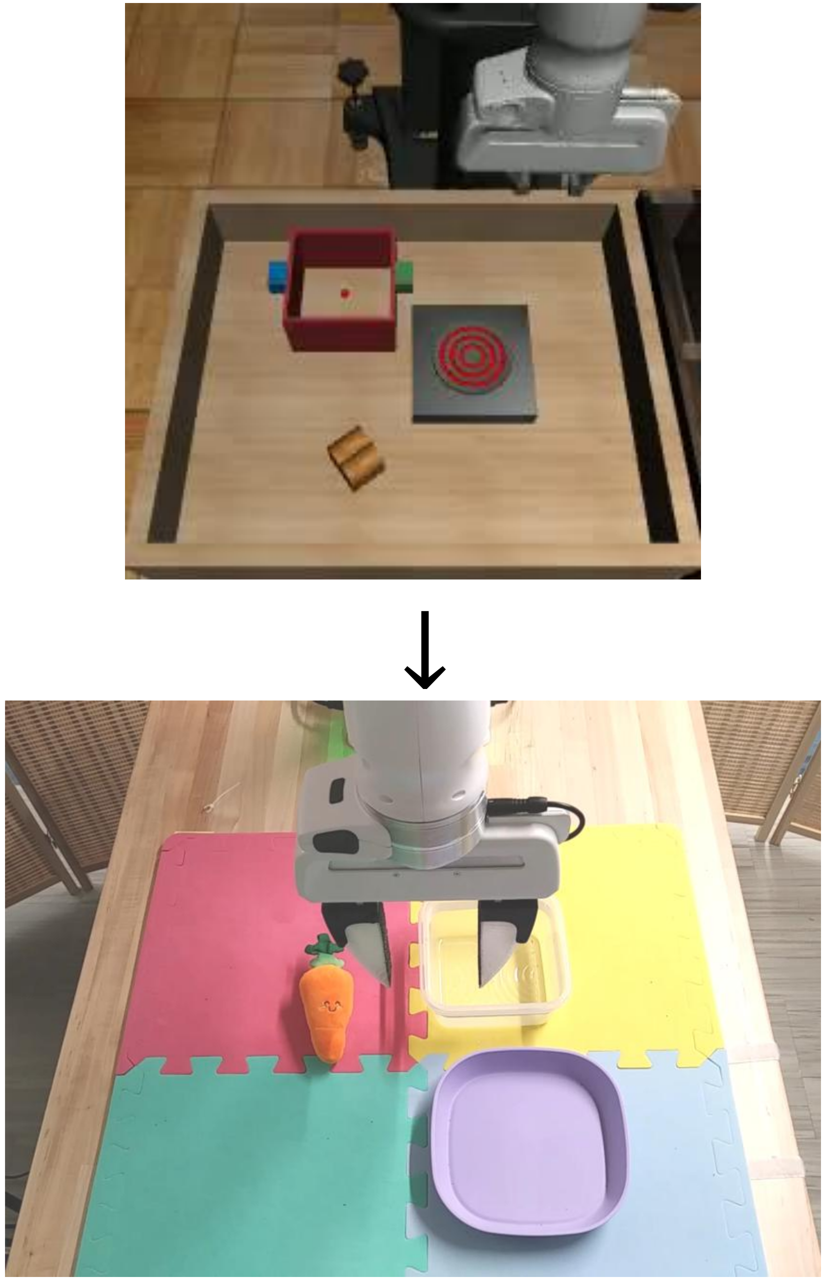

Wrap Wire

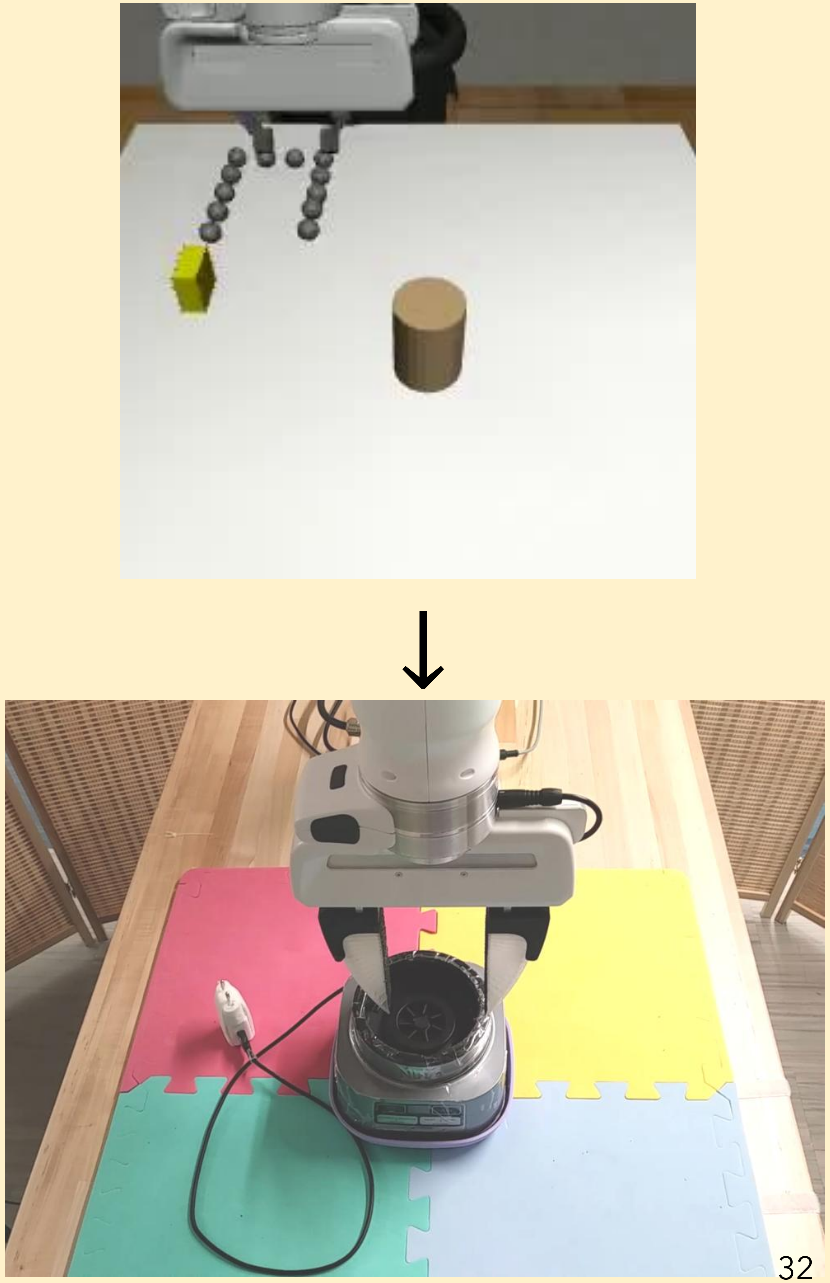

Sim2Real Results by Task: 100 real trajs.

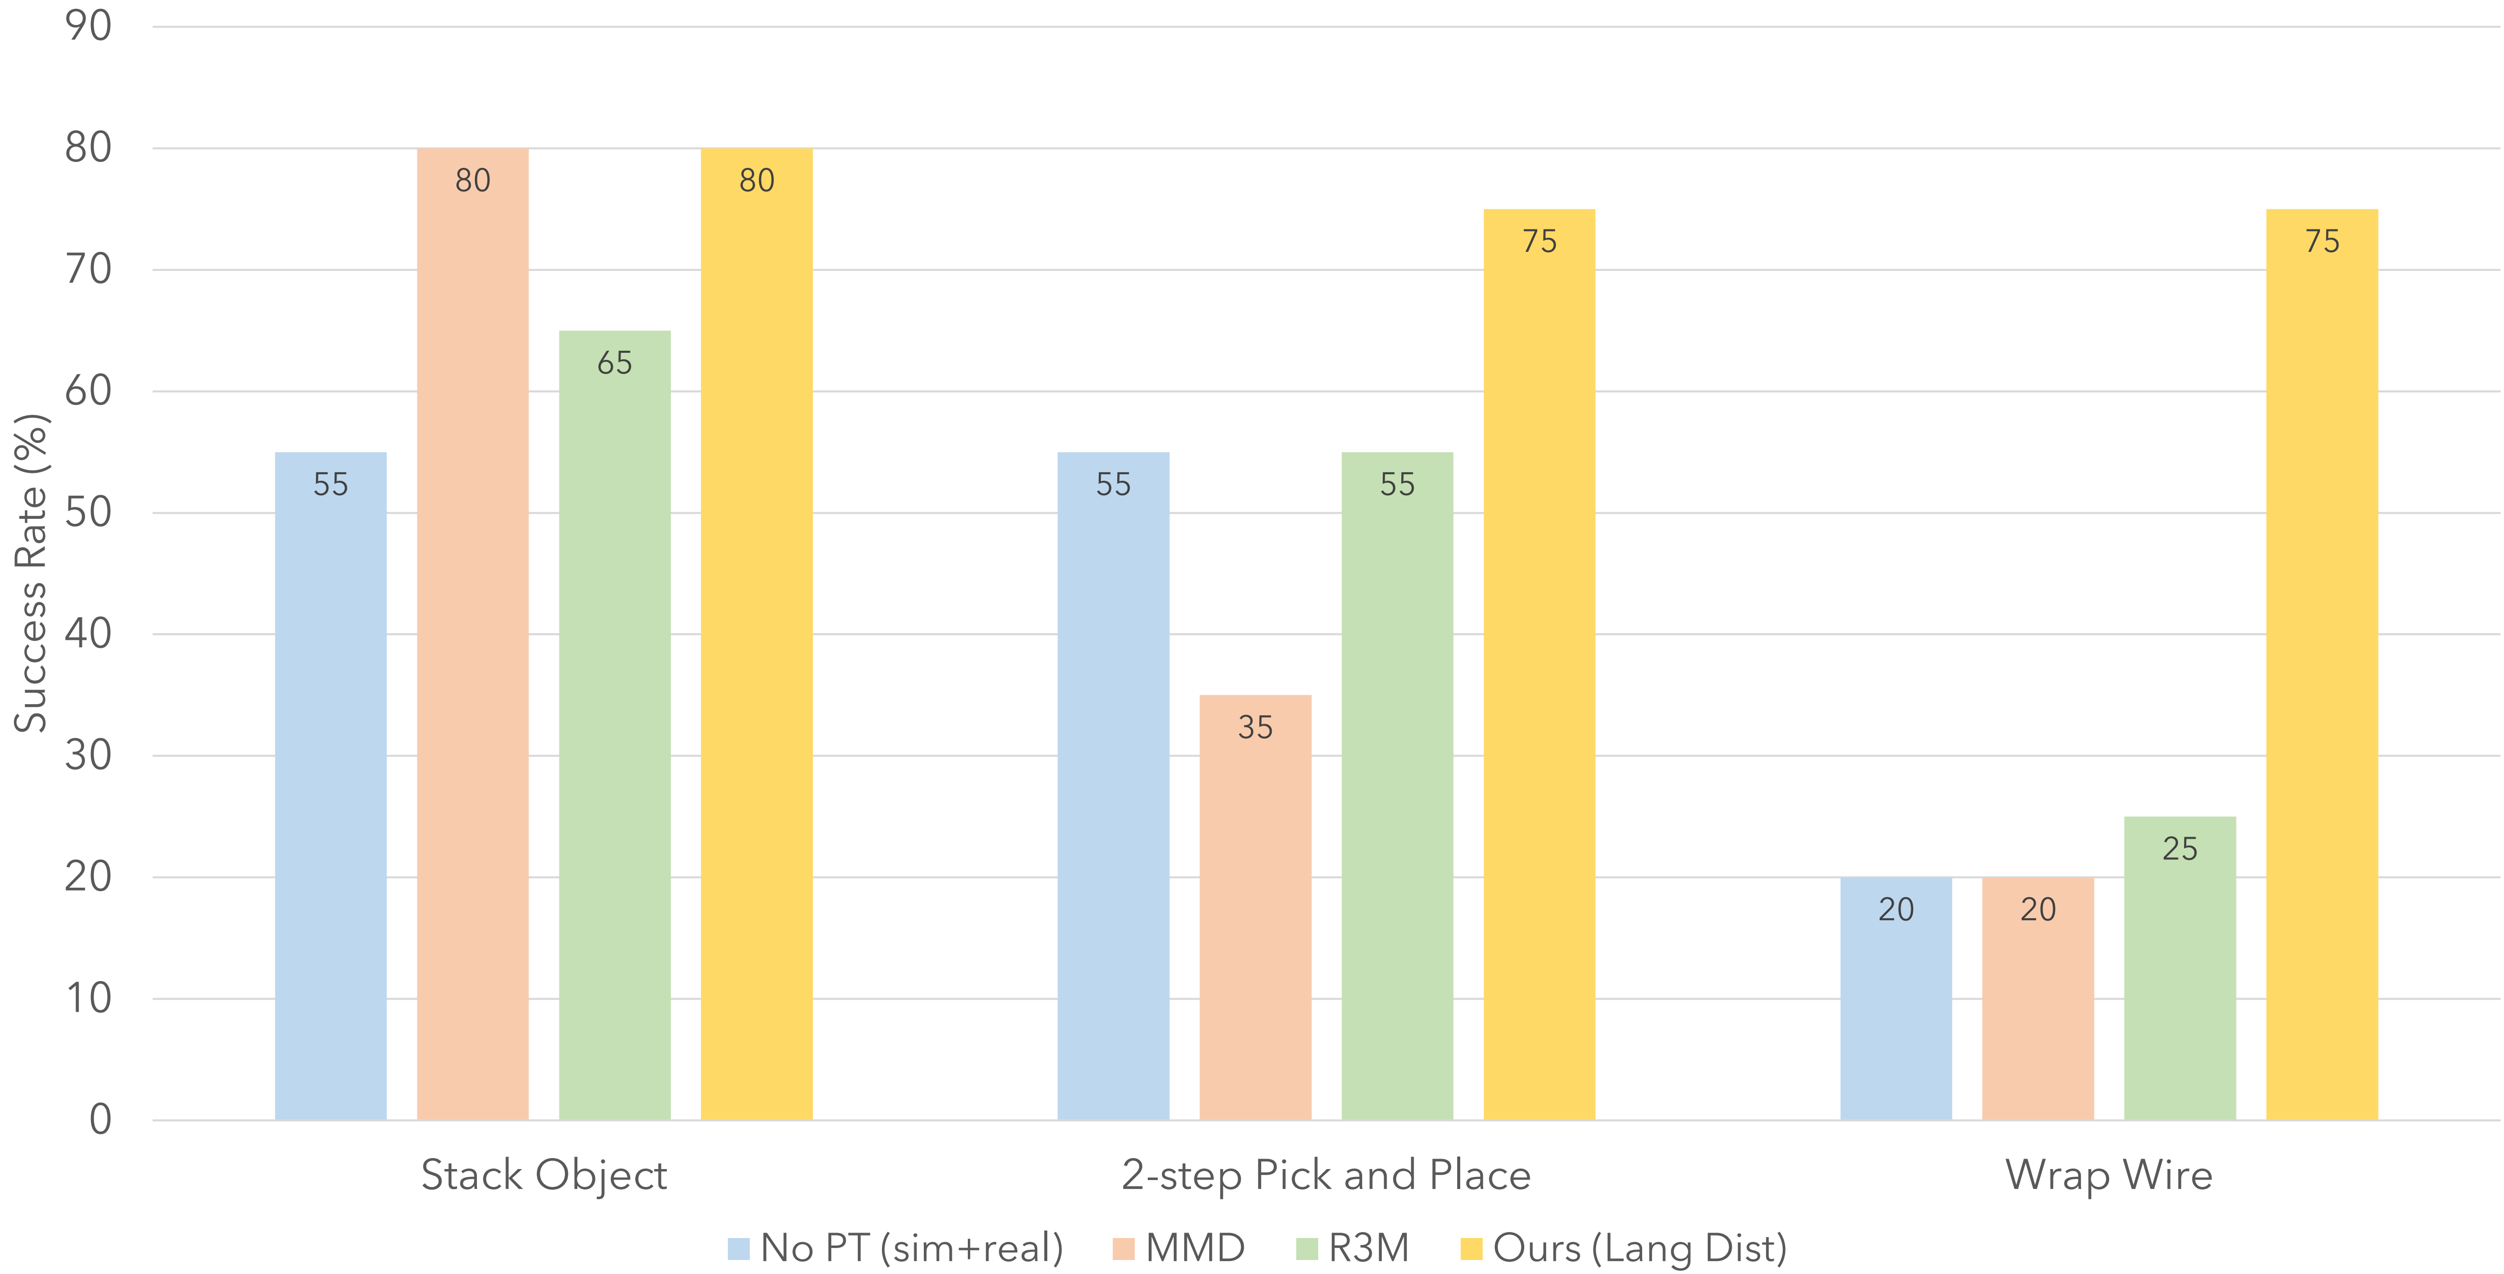

Each bar: 2 seeds, 10 trials/seed

Sim2Real Results by Task: 100 real trajs.

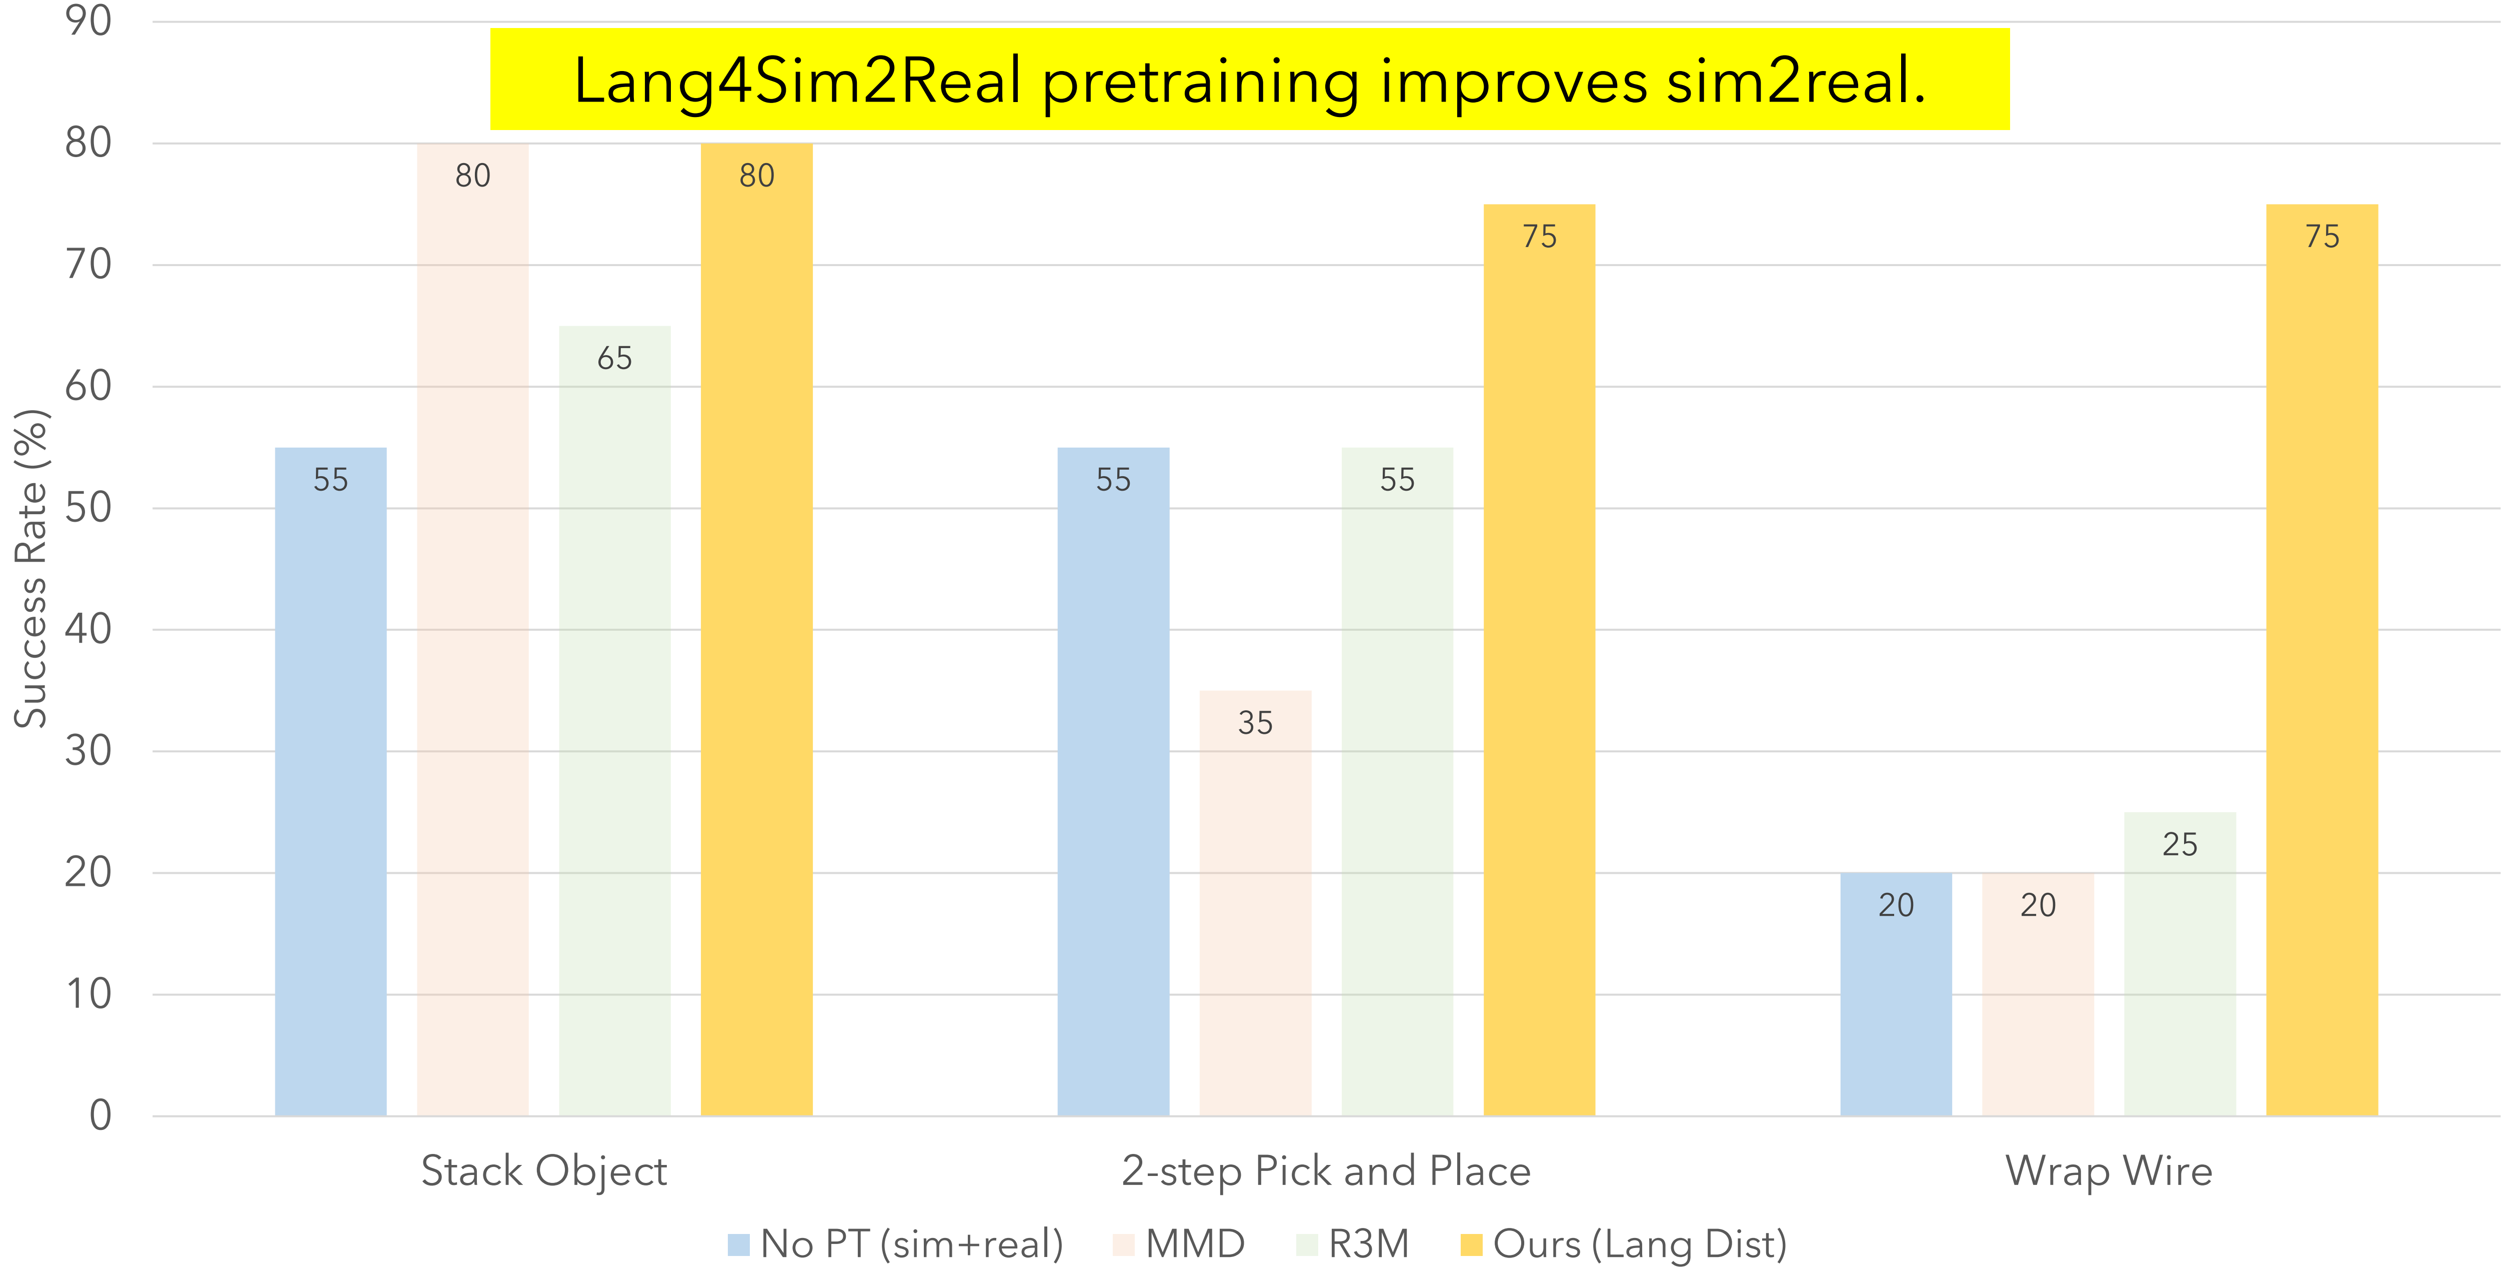

Each bar: 2 seeds, 10 trials/seed

Sim2Real Results by Task: 100 real trajs.

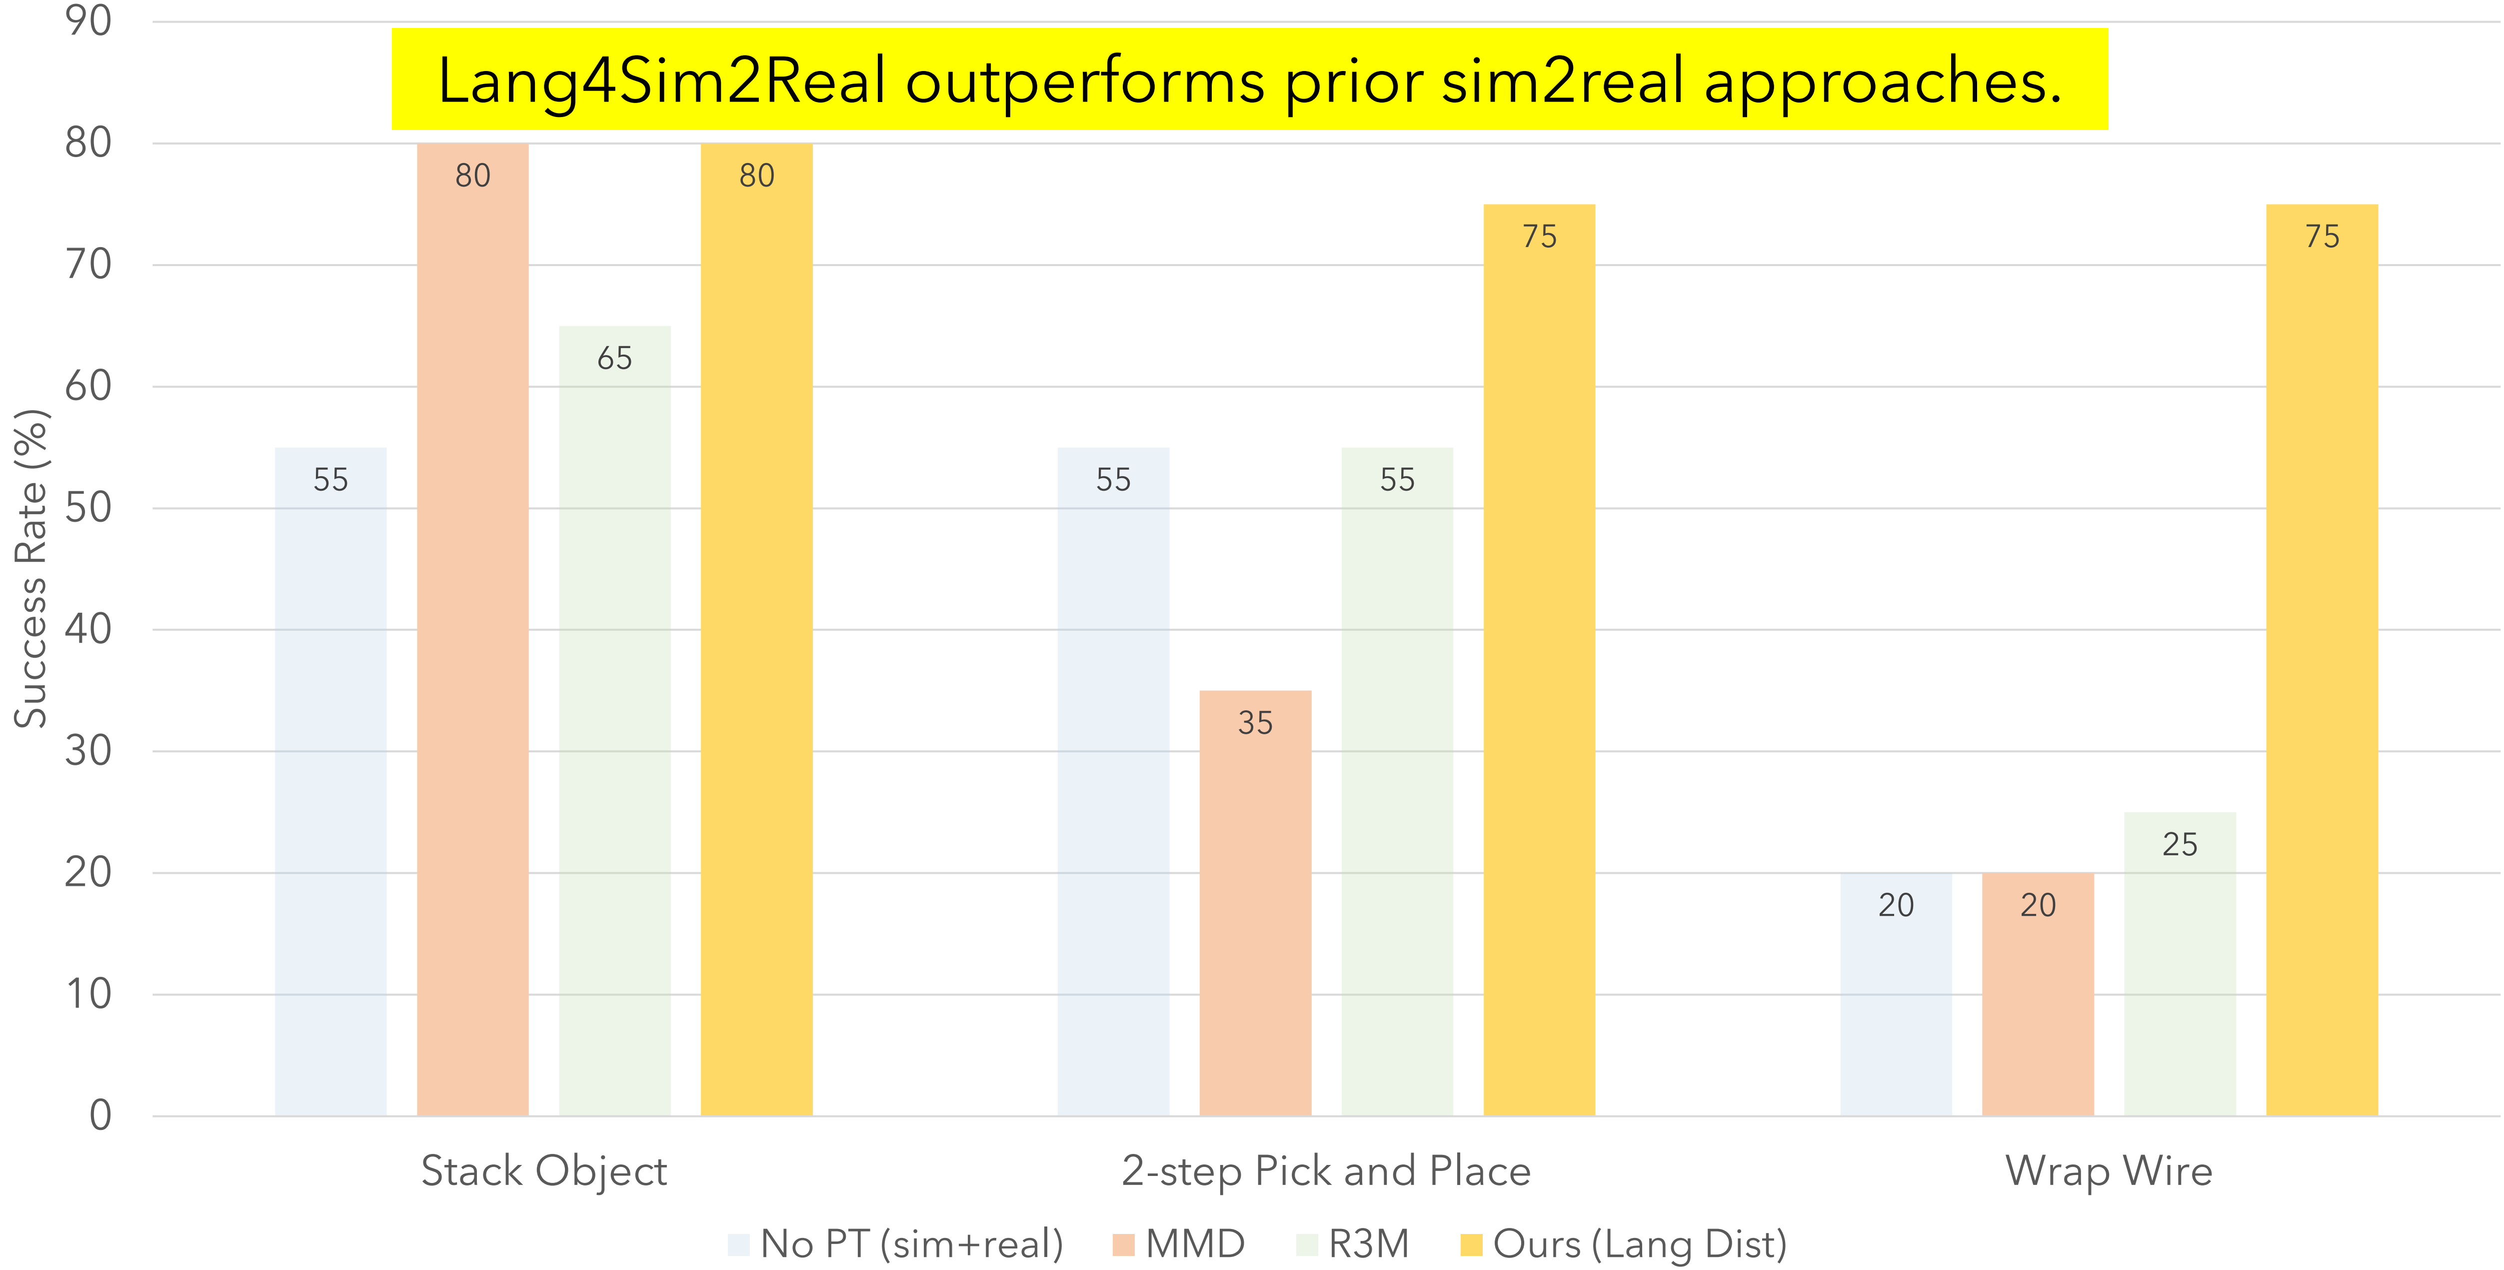

Each bar: 2 seeds, 10 trials/seed

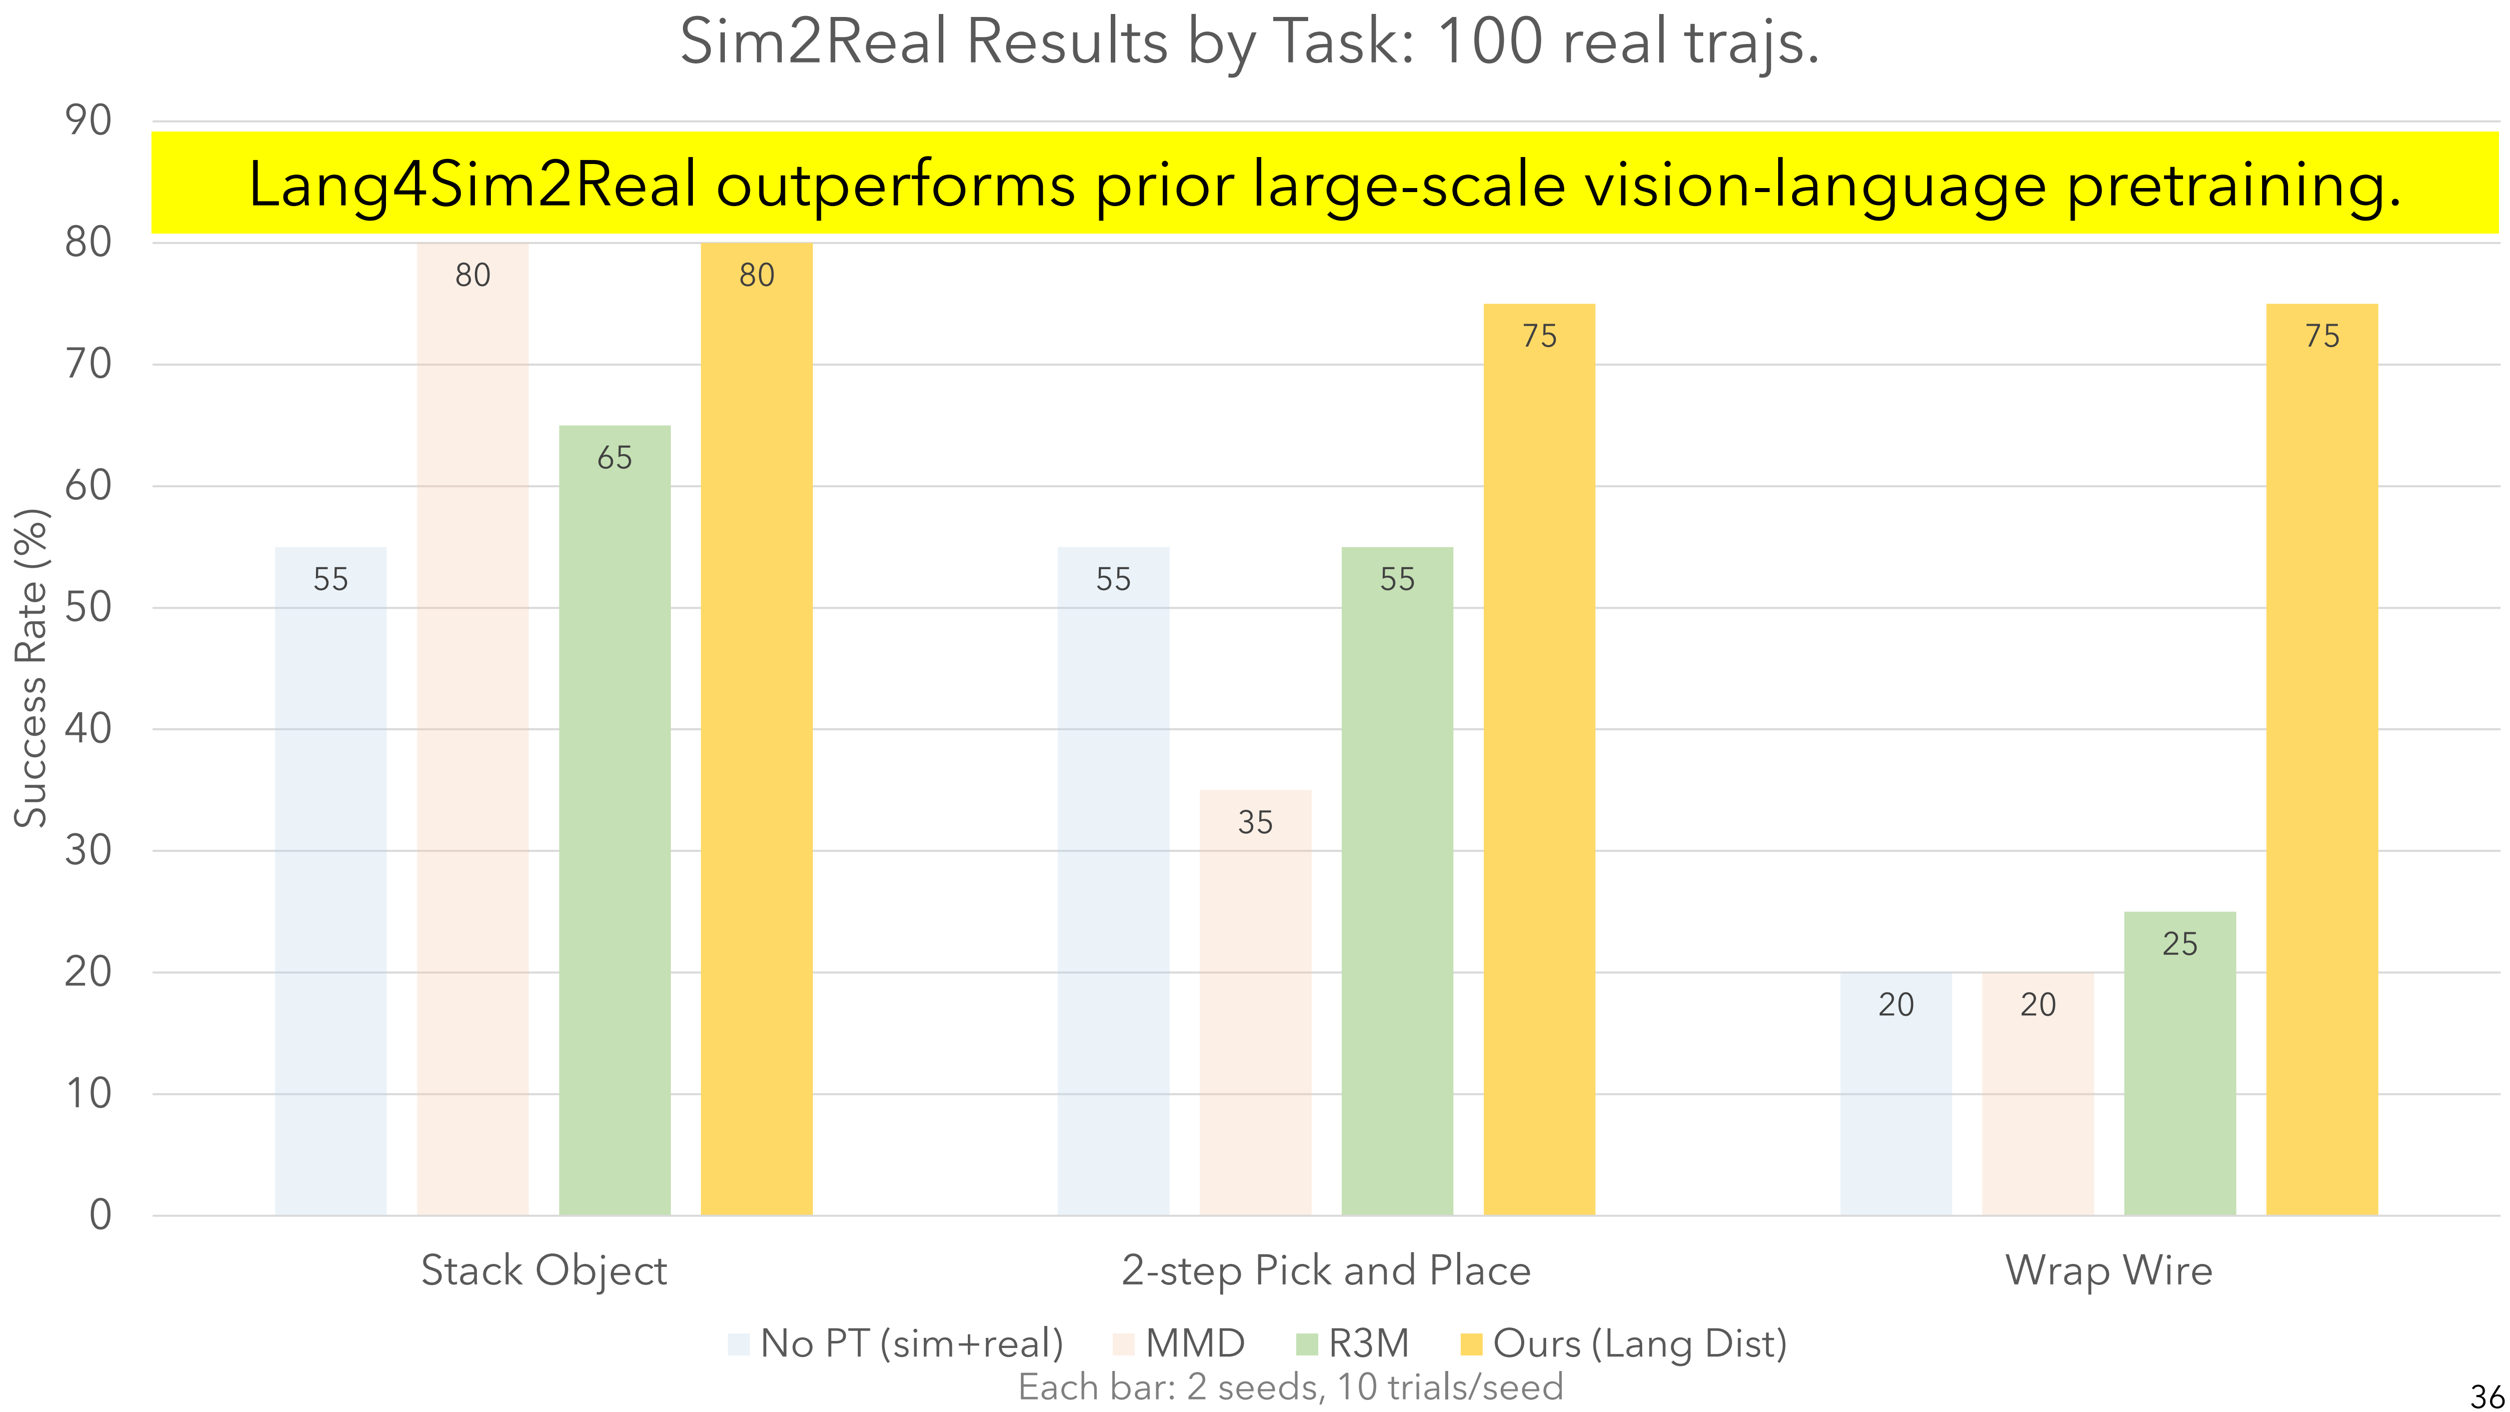

Sim2Real Results by Task: 100 real trajs.
Lang4Sim2Real outperforms prior large-scale vision-language pretraining.
90
80
70
60
50
40
30
20
10
0
Success Rate (%)
55
80
65
80
55
35
55
75
20
20
25
75
Stack Object
2-step Pick and Place
Wrap Wire
No PT (sim+real)
MMD
R3M
Ours (Lang Dist)
Each bar: 2 seeds, 10 trials/seed

Sim2Real Results by Task: 100 real trajs.

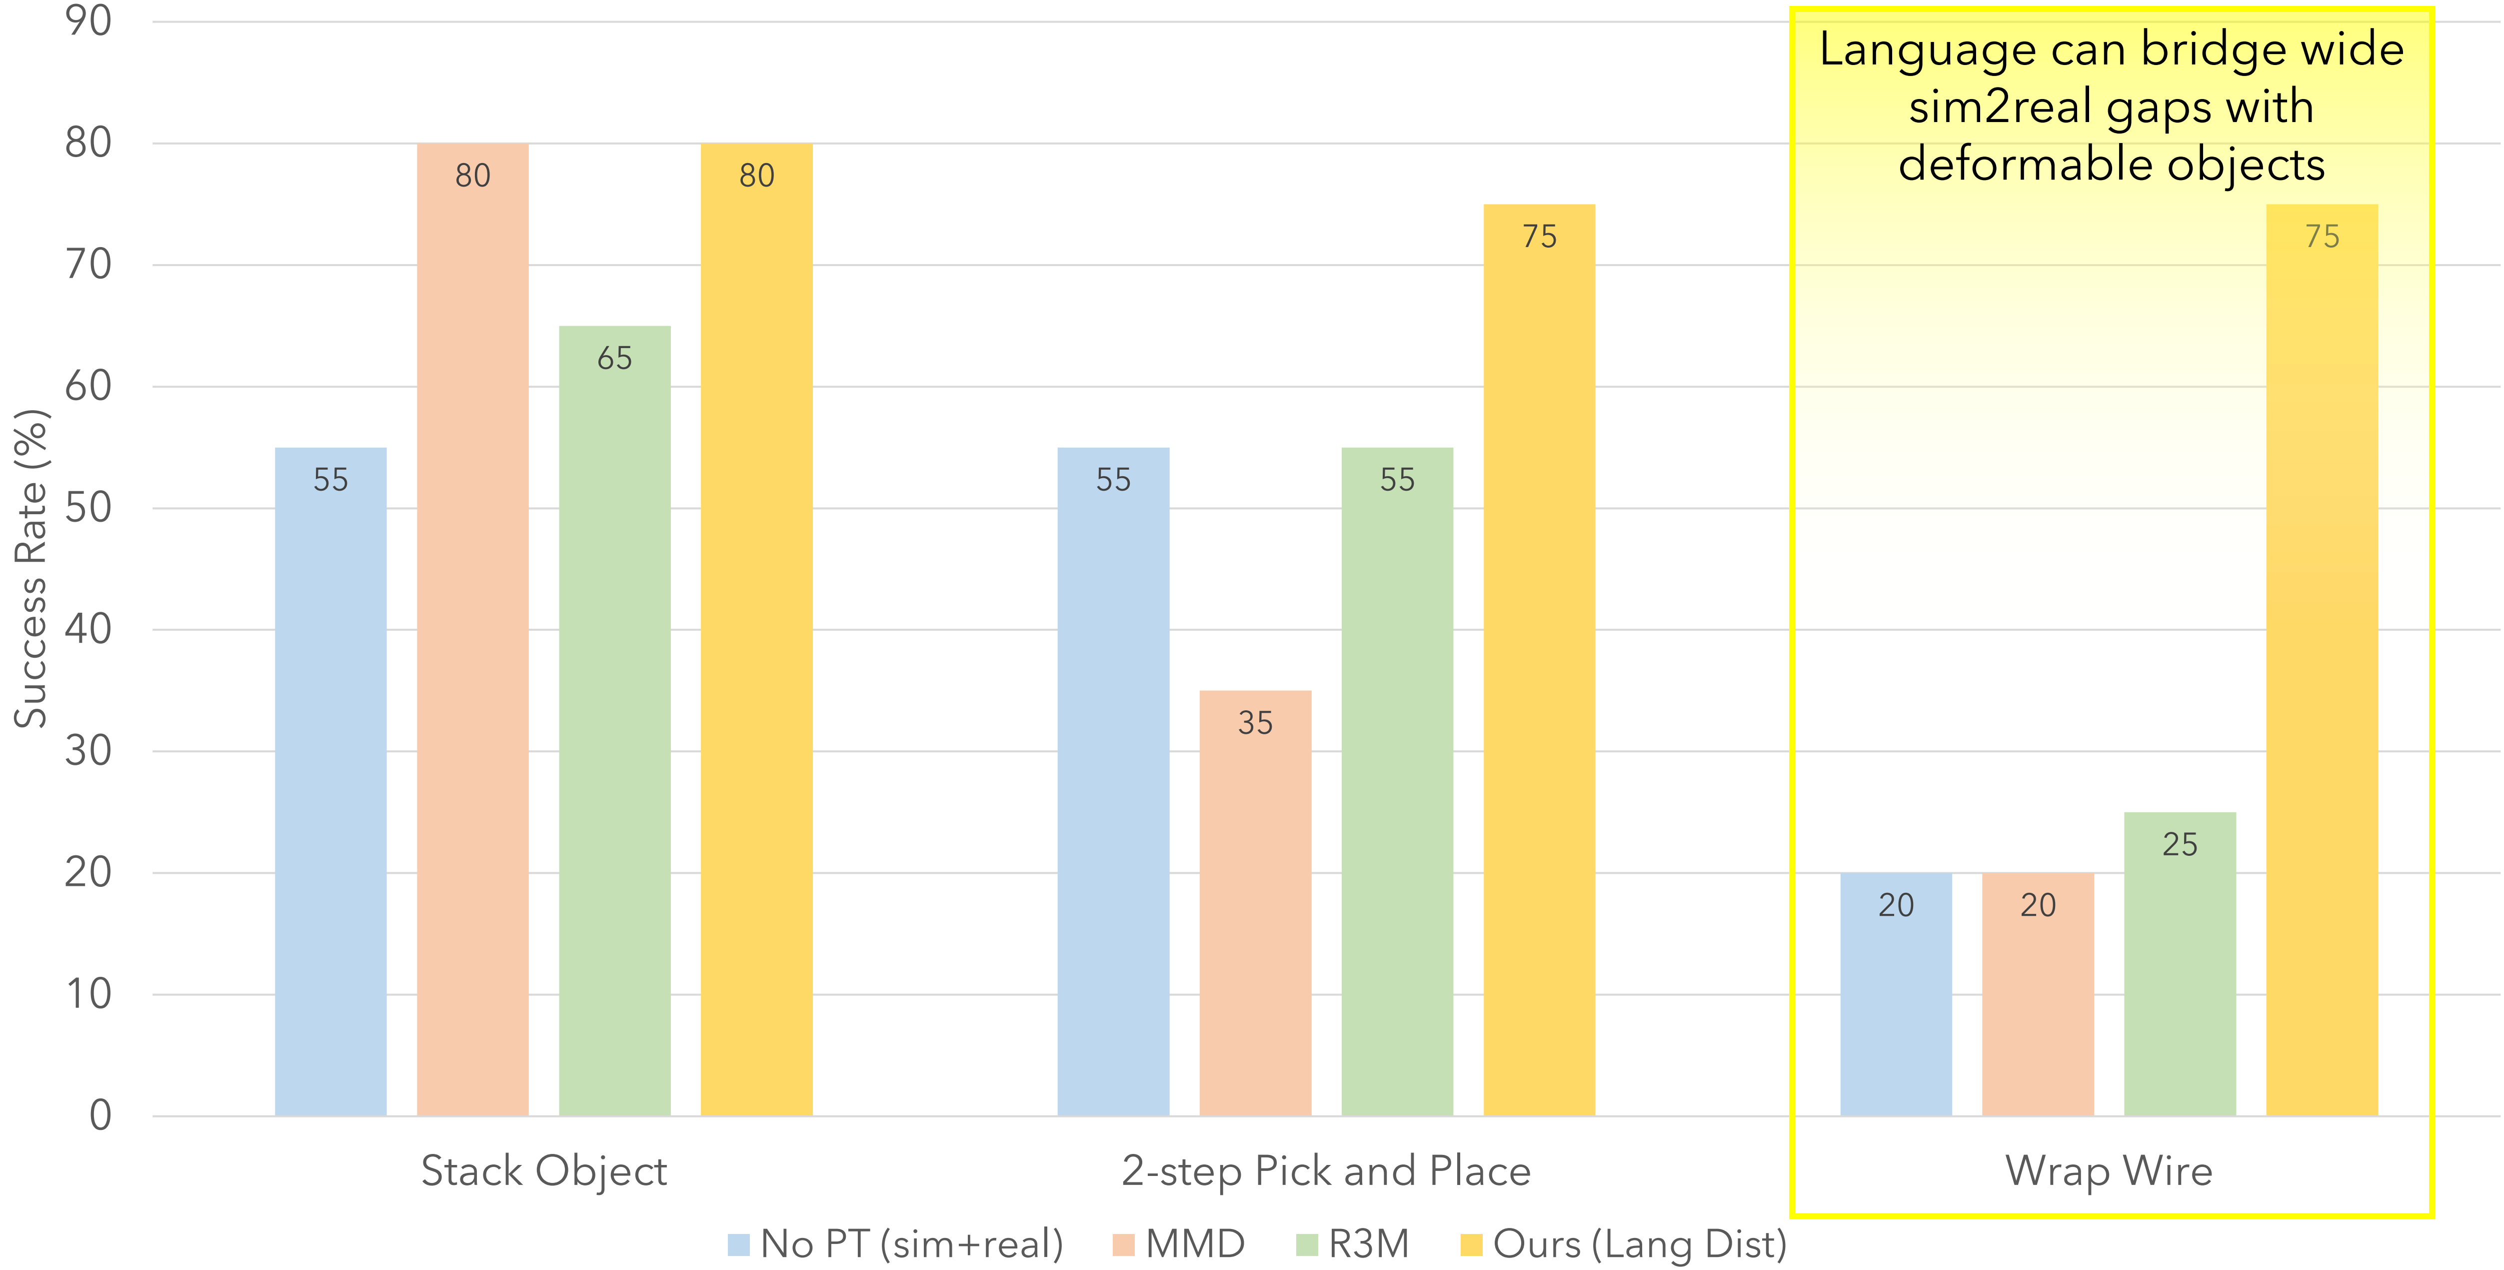

Each bar: 2 seeds, 10 trials/seed

2-step Pick and Place

Our π rollouts (5×)

Successes

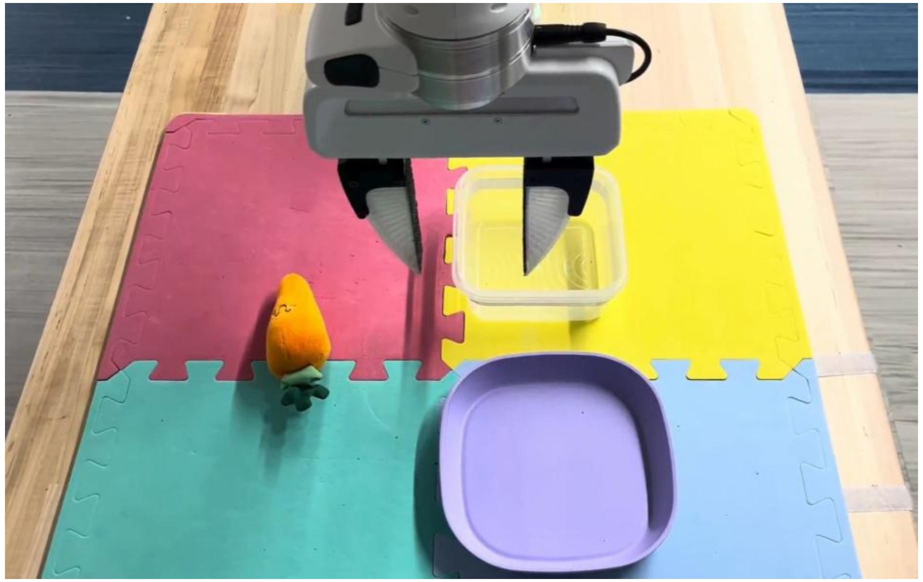

Failures

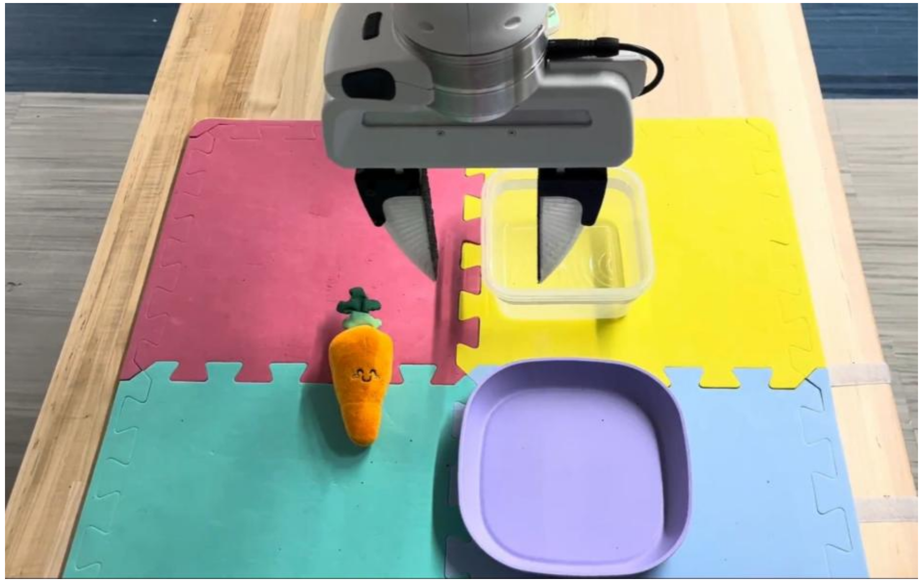

Wrap Wire

Our π rollouts (5×)

Successes

Failures

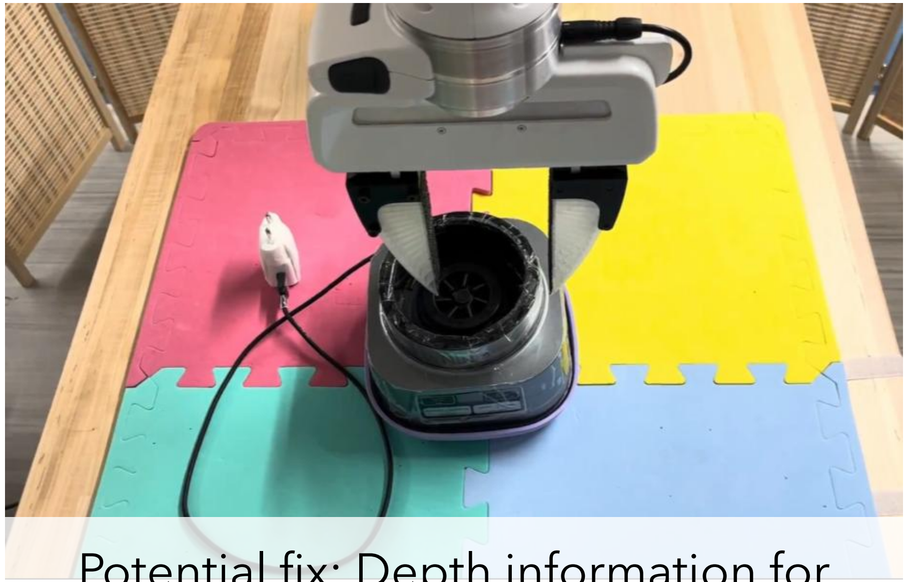

Potential fix: Depth information for more accurate grasping

Key Takeaways

Language can bridge the sim2real gap with domain-invariant representations.

Our method enables leveraging low-fidelity sim data for sim2real transfer **on deformable objects**.

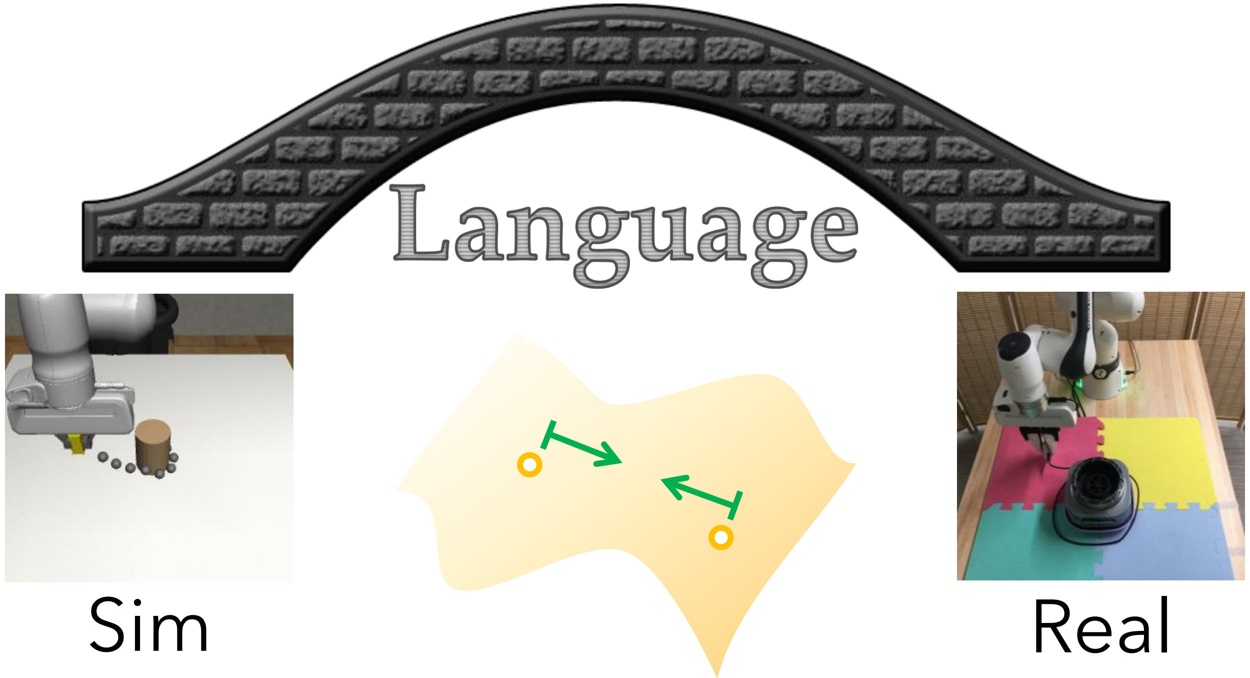

Outline

Language Improves Robotic Capabilities through…

Semantics for Generalization

1. DeL-TaCo: Simultaneous Learning from Demo + Language
2. Lang4Sim2Real: Language to Bridge the Sim2Real Gap

Communicative Medium

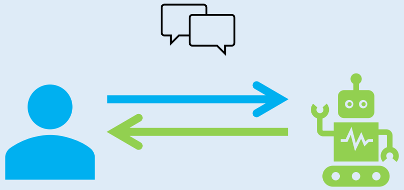

3. **MicoBot**: Bidirectional Dialog for Human-Robot Collaboration
4. BookWorm: Learning to operate new devices with guidebooks and human feedback

Future Work and Conclusion

Mixed-Initiative Dialog for Human-Robot Collaborative Manipulation

Collaborators: Roberto Martín-Martín, Ray Mooney, Chengshu Li, Luca Macesanu, Arnav Balaji, Ruchira Ray

ICRA 2026

Imagine preparing for a dinner party with a friend…


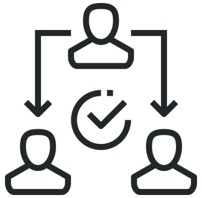
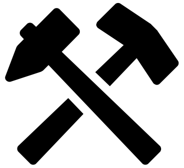


Allocate tasks by skillset

I can cook, but don’t know how to make drinks. Can you?

Sure, I make the best smoothies.



Can you help wash dishes?

No, I’m busy vacuuming the floor.


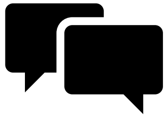


Adapt to availability and willingness

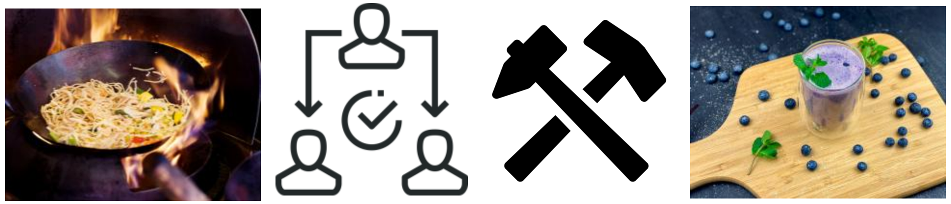

Allocate tasks by skillset

I can cook, but don't know how to make drinks. Can you?

Sure, I make the best smoothies.

Both initiate dialog to collaborate better

Can you help wash dishes?

No, I'm busy vacuuming the floor.

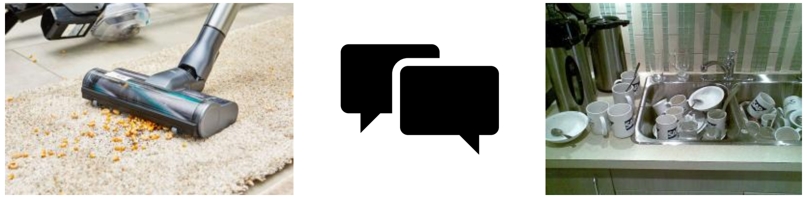

Adapt to availability and willingness

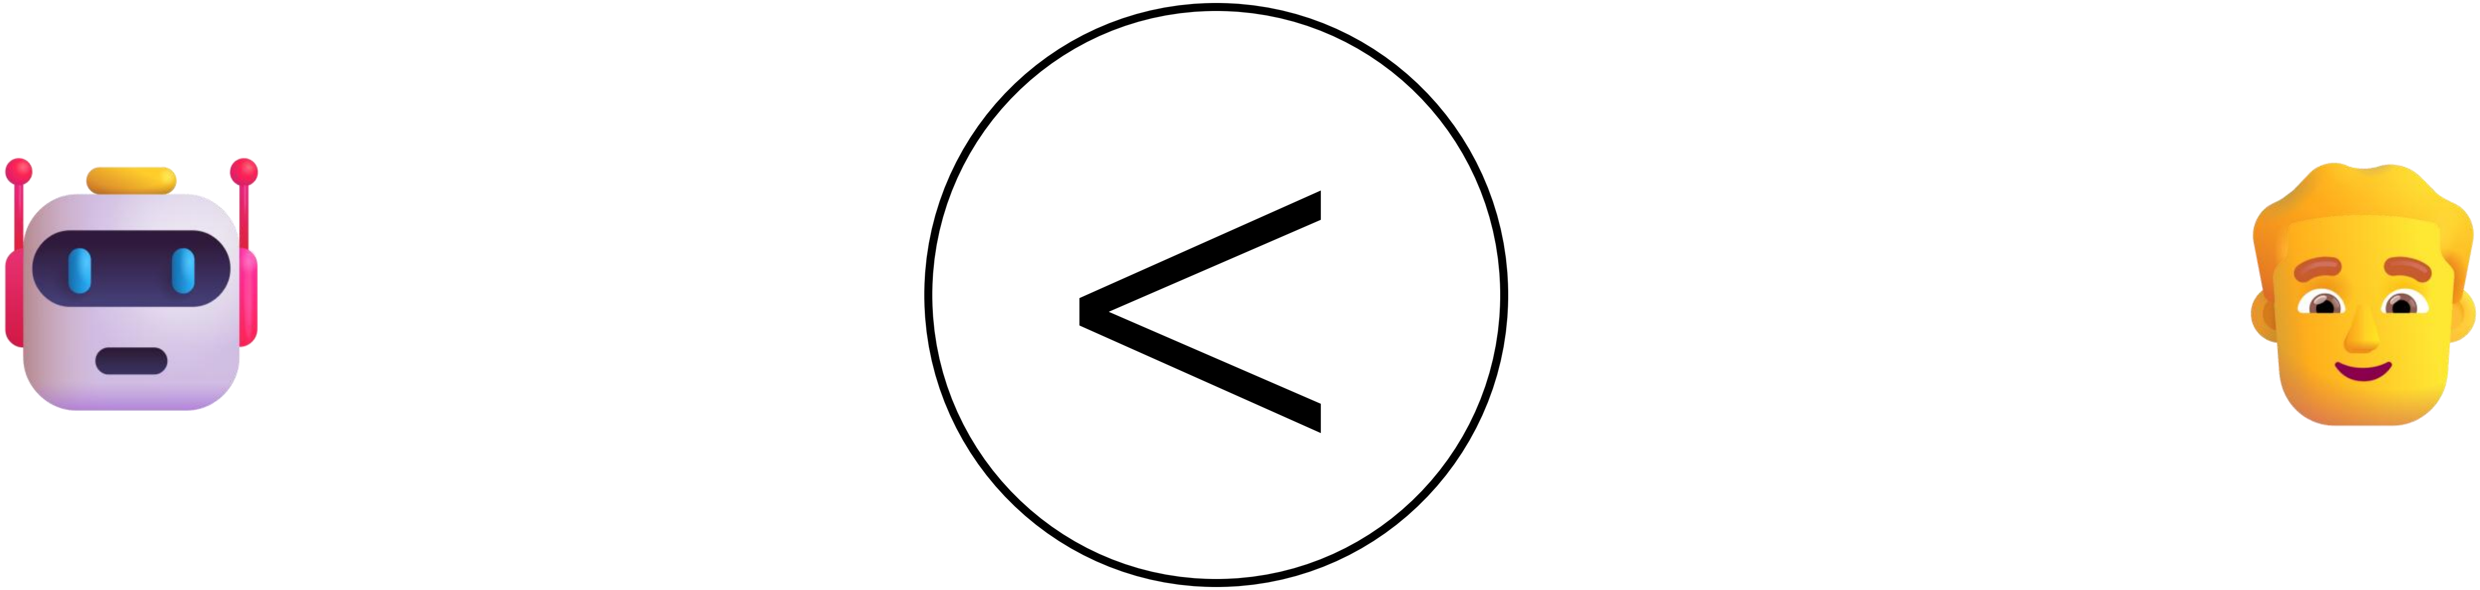

Robots struggle on complex tasks…

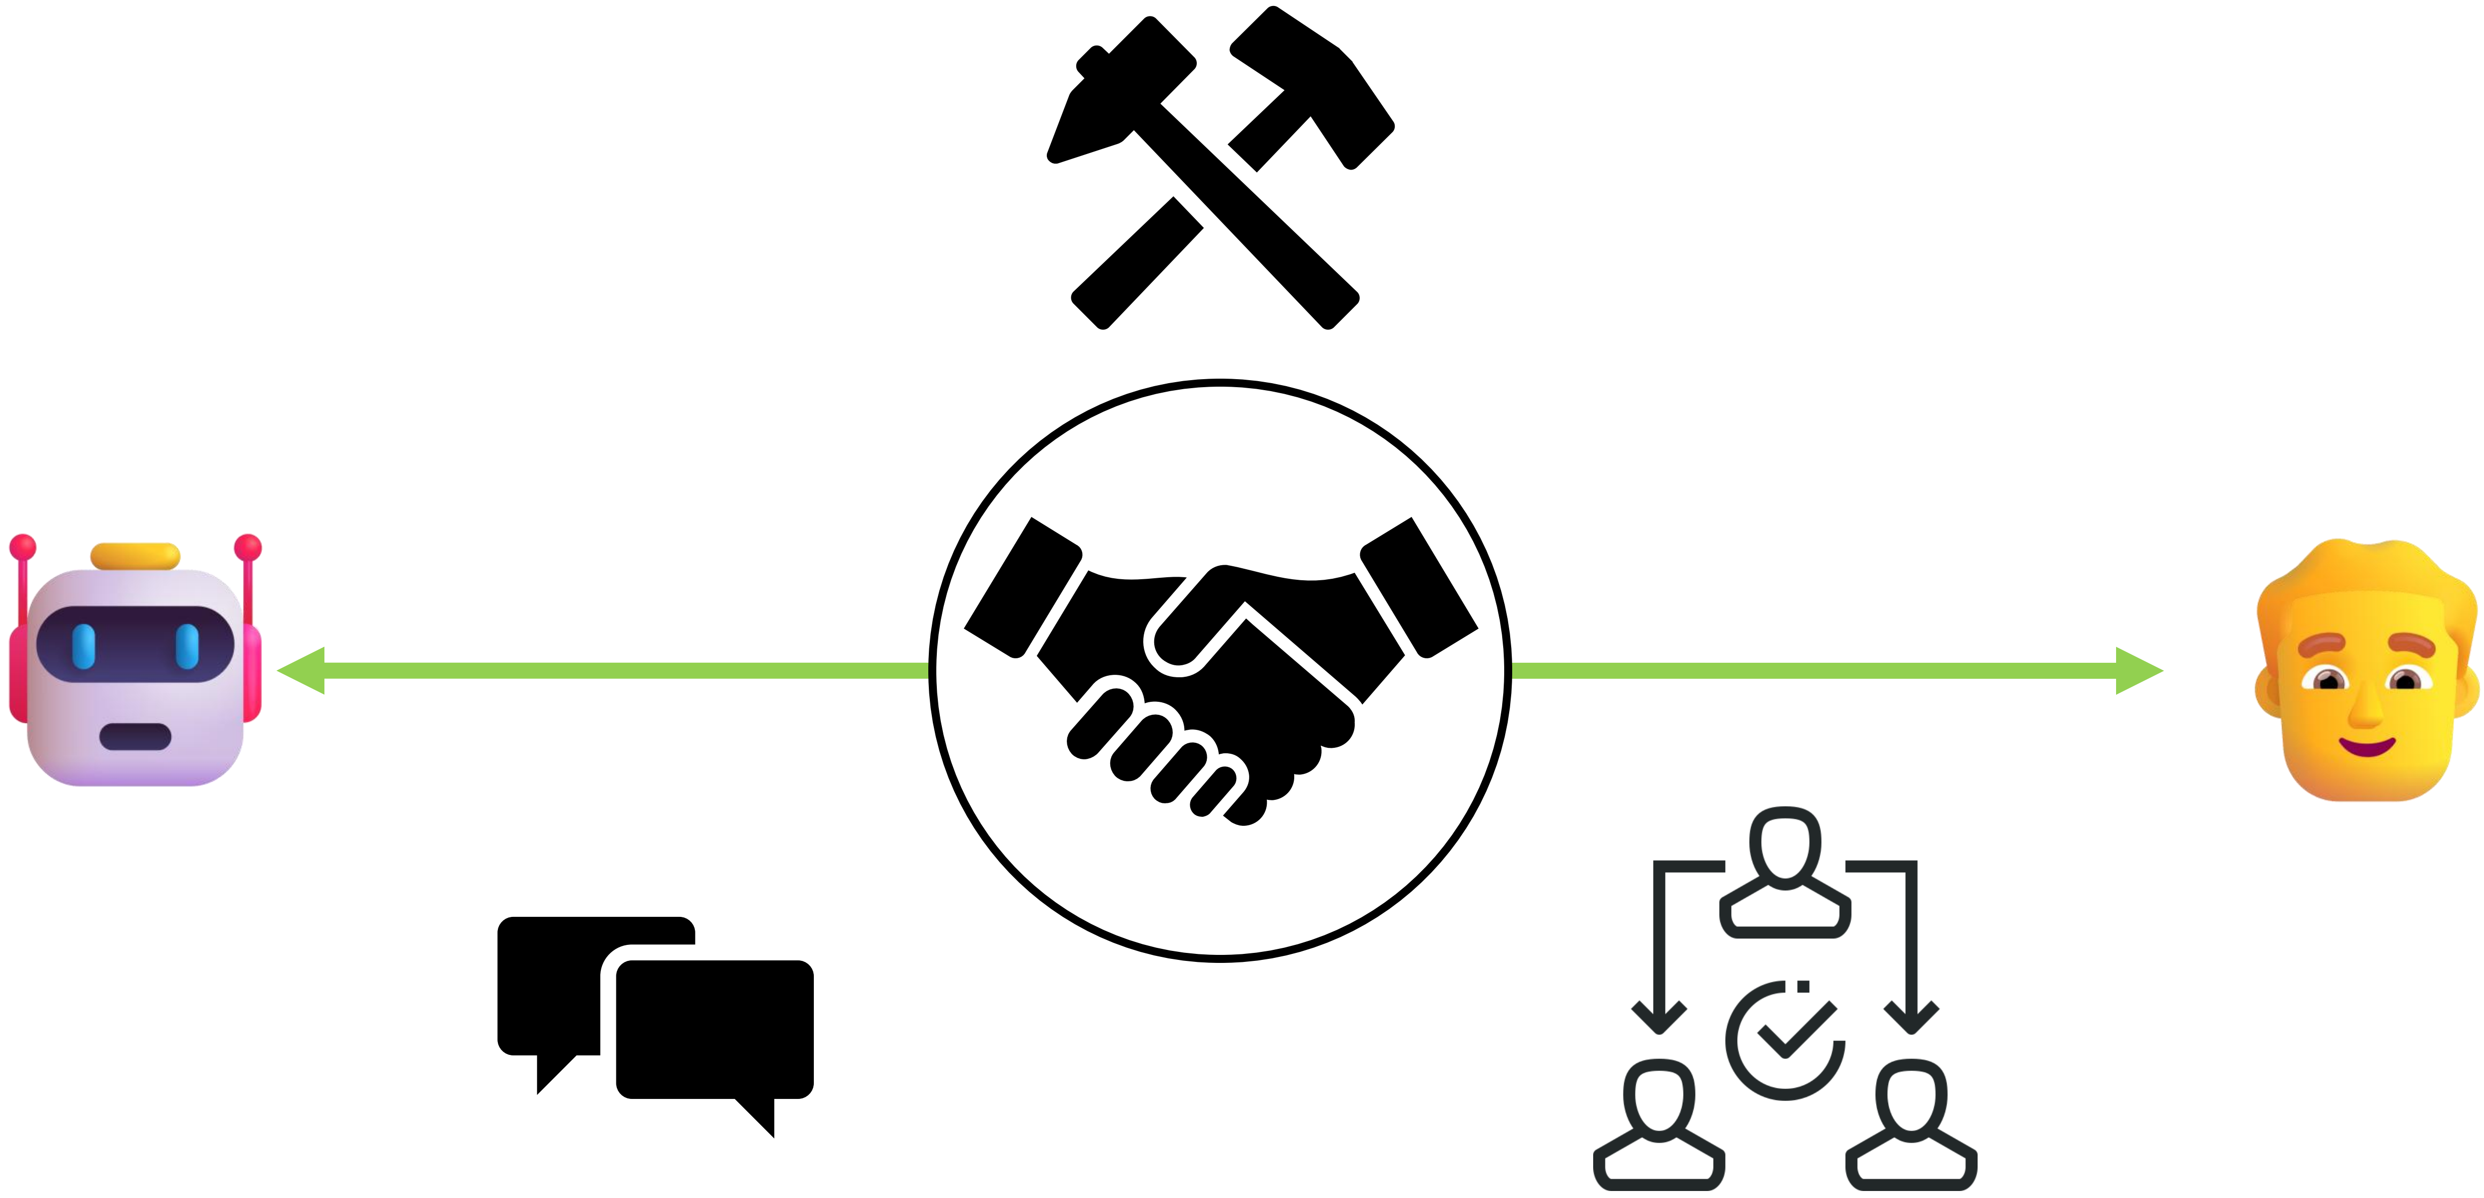

…and need to collaborate and communicate with humans.

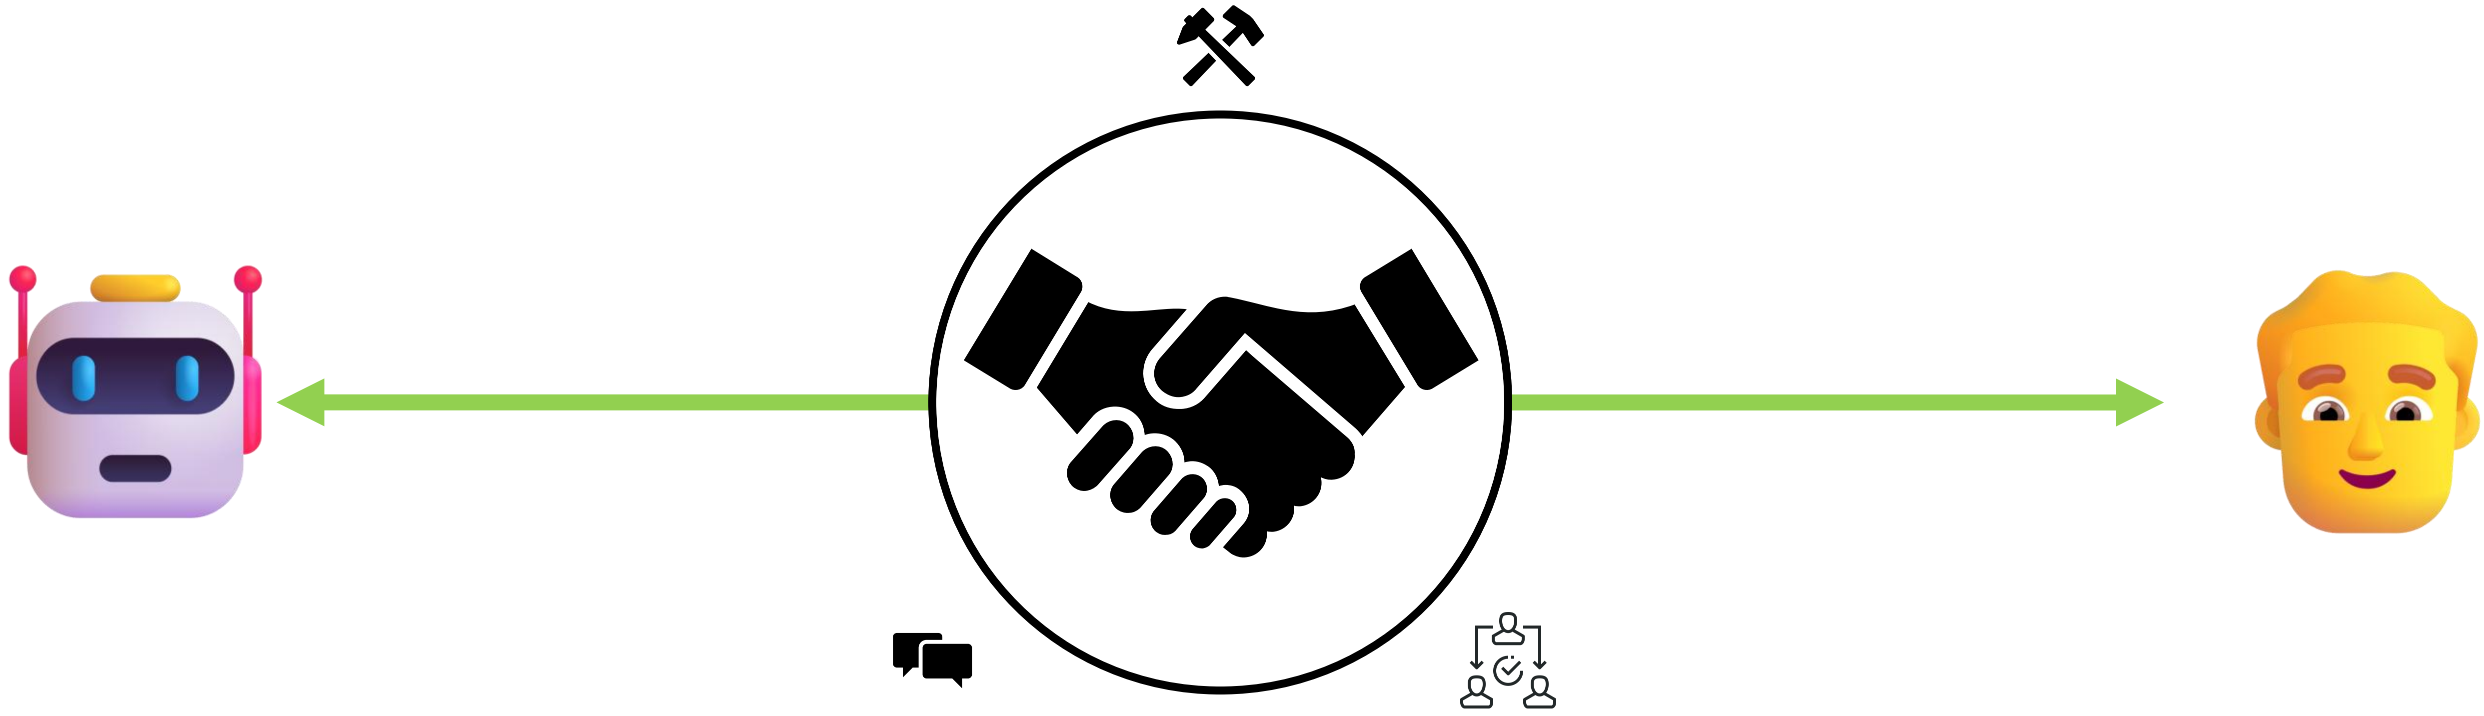

Hypothesis:

A collaborative robot must leverage 🤖 + 👨 physical capabilities and adapt to changing 👨 preferences and helpfulness through **bidirectional communication**.

Prior Work

Human-Initiated Dialog
Most Modern Chatbots

Robot-Initiated Dialog
In Human-Robot Interaction (HRI)
MOSAIC (Wang+, CoRL 2024)
RoCo (Mandi+, ICRA 2024)
KnowNo (Ren+, CoRL 2023)

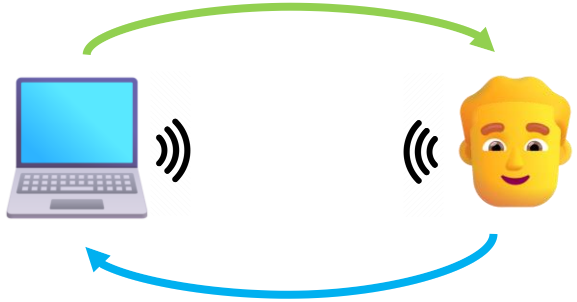

Mixed-Initiative Dialog
Jaime R. Carbonell, 1970
IEEE Man-Machine Systems
(Deng+, 2023)
(Chen+, ACL 2023)

Single initiative: limited adaptability for collaboration | Not shown on physical robots

MOSAIC: Huaxiaoyue Wang, Kushal Kedia, Juntao Ren, Rahma Abdullah, Atiksh Bhardwaj, Angela Chao, Kelly Y Chen, Nathaniel Chin, Prithwish Dan, Xinyi Fan, Gonzalo Gonzalez-Pumariega, Aditya Kompella, Maximus Adrian Pace, Yash Sharma, Xiangwan Sun, Neha Sunkara, Sanjiban Choudhury. "MOSAIC: Modular Foundation Models for Assistive and Interactive Cooking." CoRL 2024.
RoCo: Zhao Mandi, Shreeya Jain, Shuran Song. "RoCo: Dialectic Multi-Robot Collaboration with LLMs." ICRA 2024.
KnowNo: Allen Z. Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, Zhenjia Xu, Dorsa Sadigh, Andy Zeng, Anirudha Majumdar. "Robots That Ask For Help: Uncertainty Alignment for Large Language Model Planners." CoRL 2023.
Deng: Yang Deng, Wenxuan Zhang, Yifei Yuan, Wai Lam. "Knowledge-enhanced Mixed-initiative Dialogue System for Emotional Support Conversations." ACL 2023.
Chen: Maximillian Chen, Xiao Yu, Weiyan Shi, Urvi Awasthi, Zhou Yu. "Controllable Mixed-Initiative Dialogue Generation through Prompting." ACL 2023.

Why Mixed-Initiative Dialog?

Ask for help when it can't do something.
Negotiate task divisions.

Offer to do sensitive steps.
Say what they won't do.

MICoBot

Mixed **I**nitiative **Co**llaborative Ro**bot**

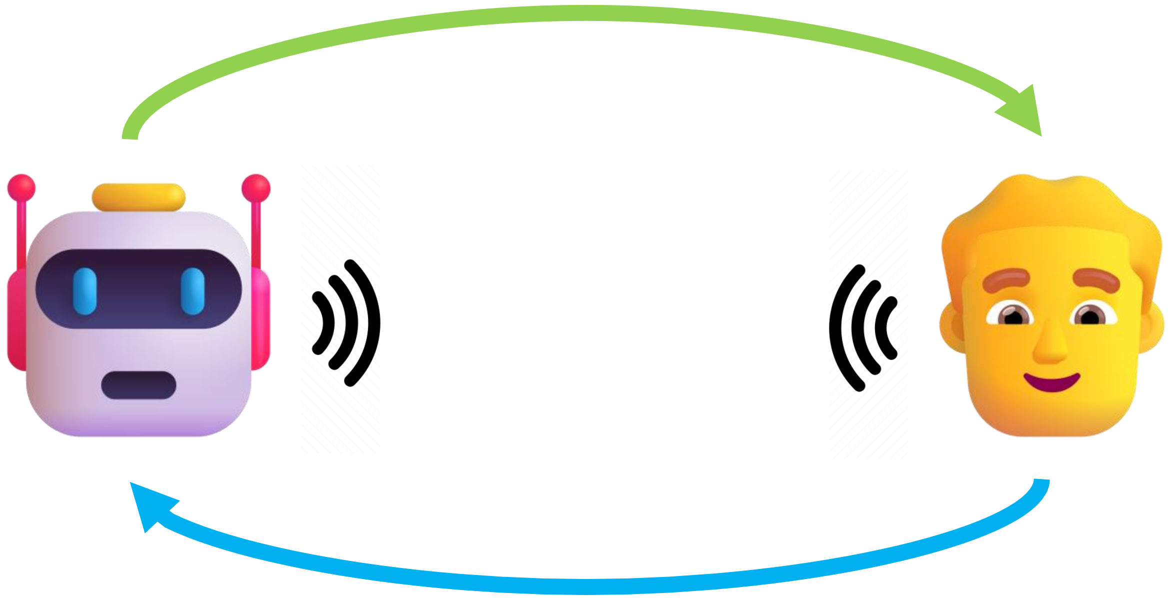

Real-world human + robot manipulation with mixed-initiative dialog

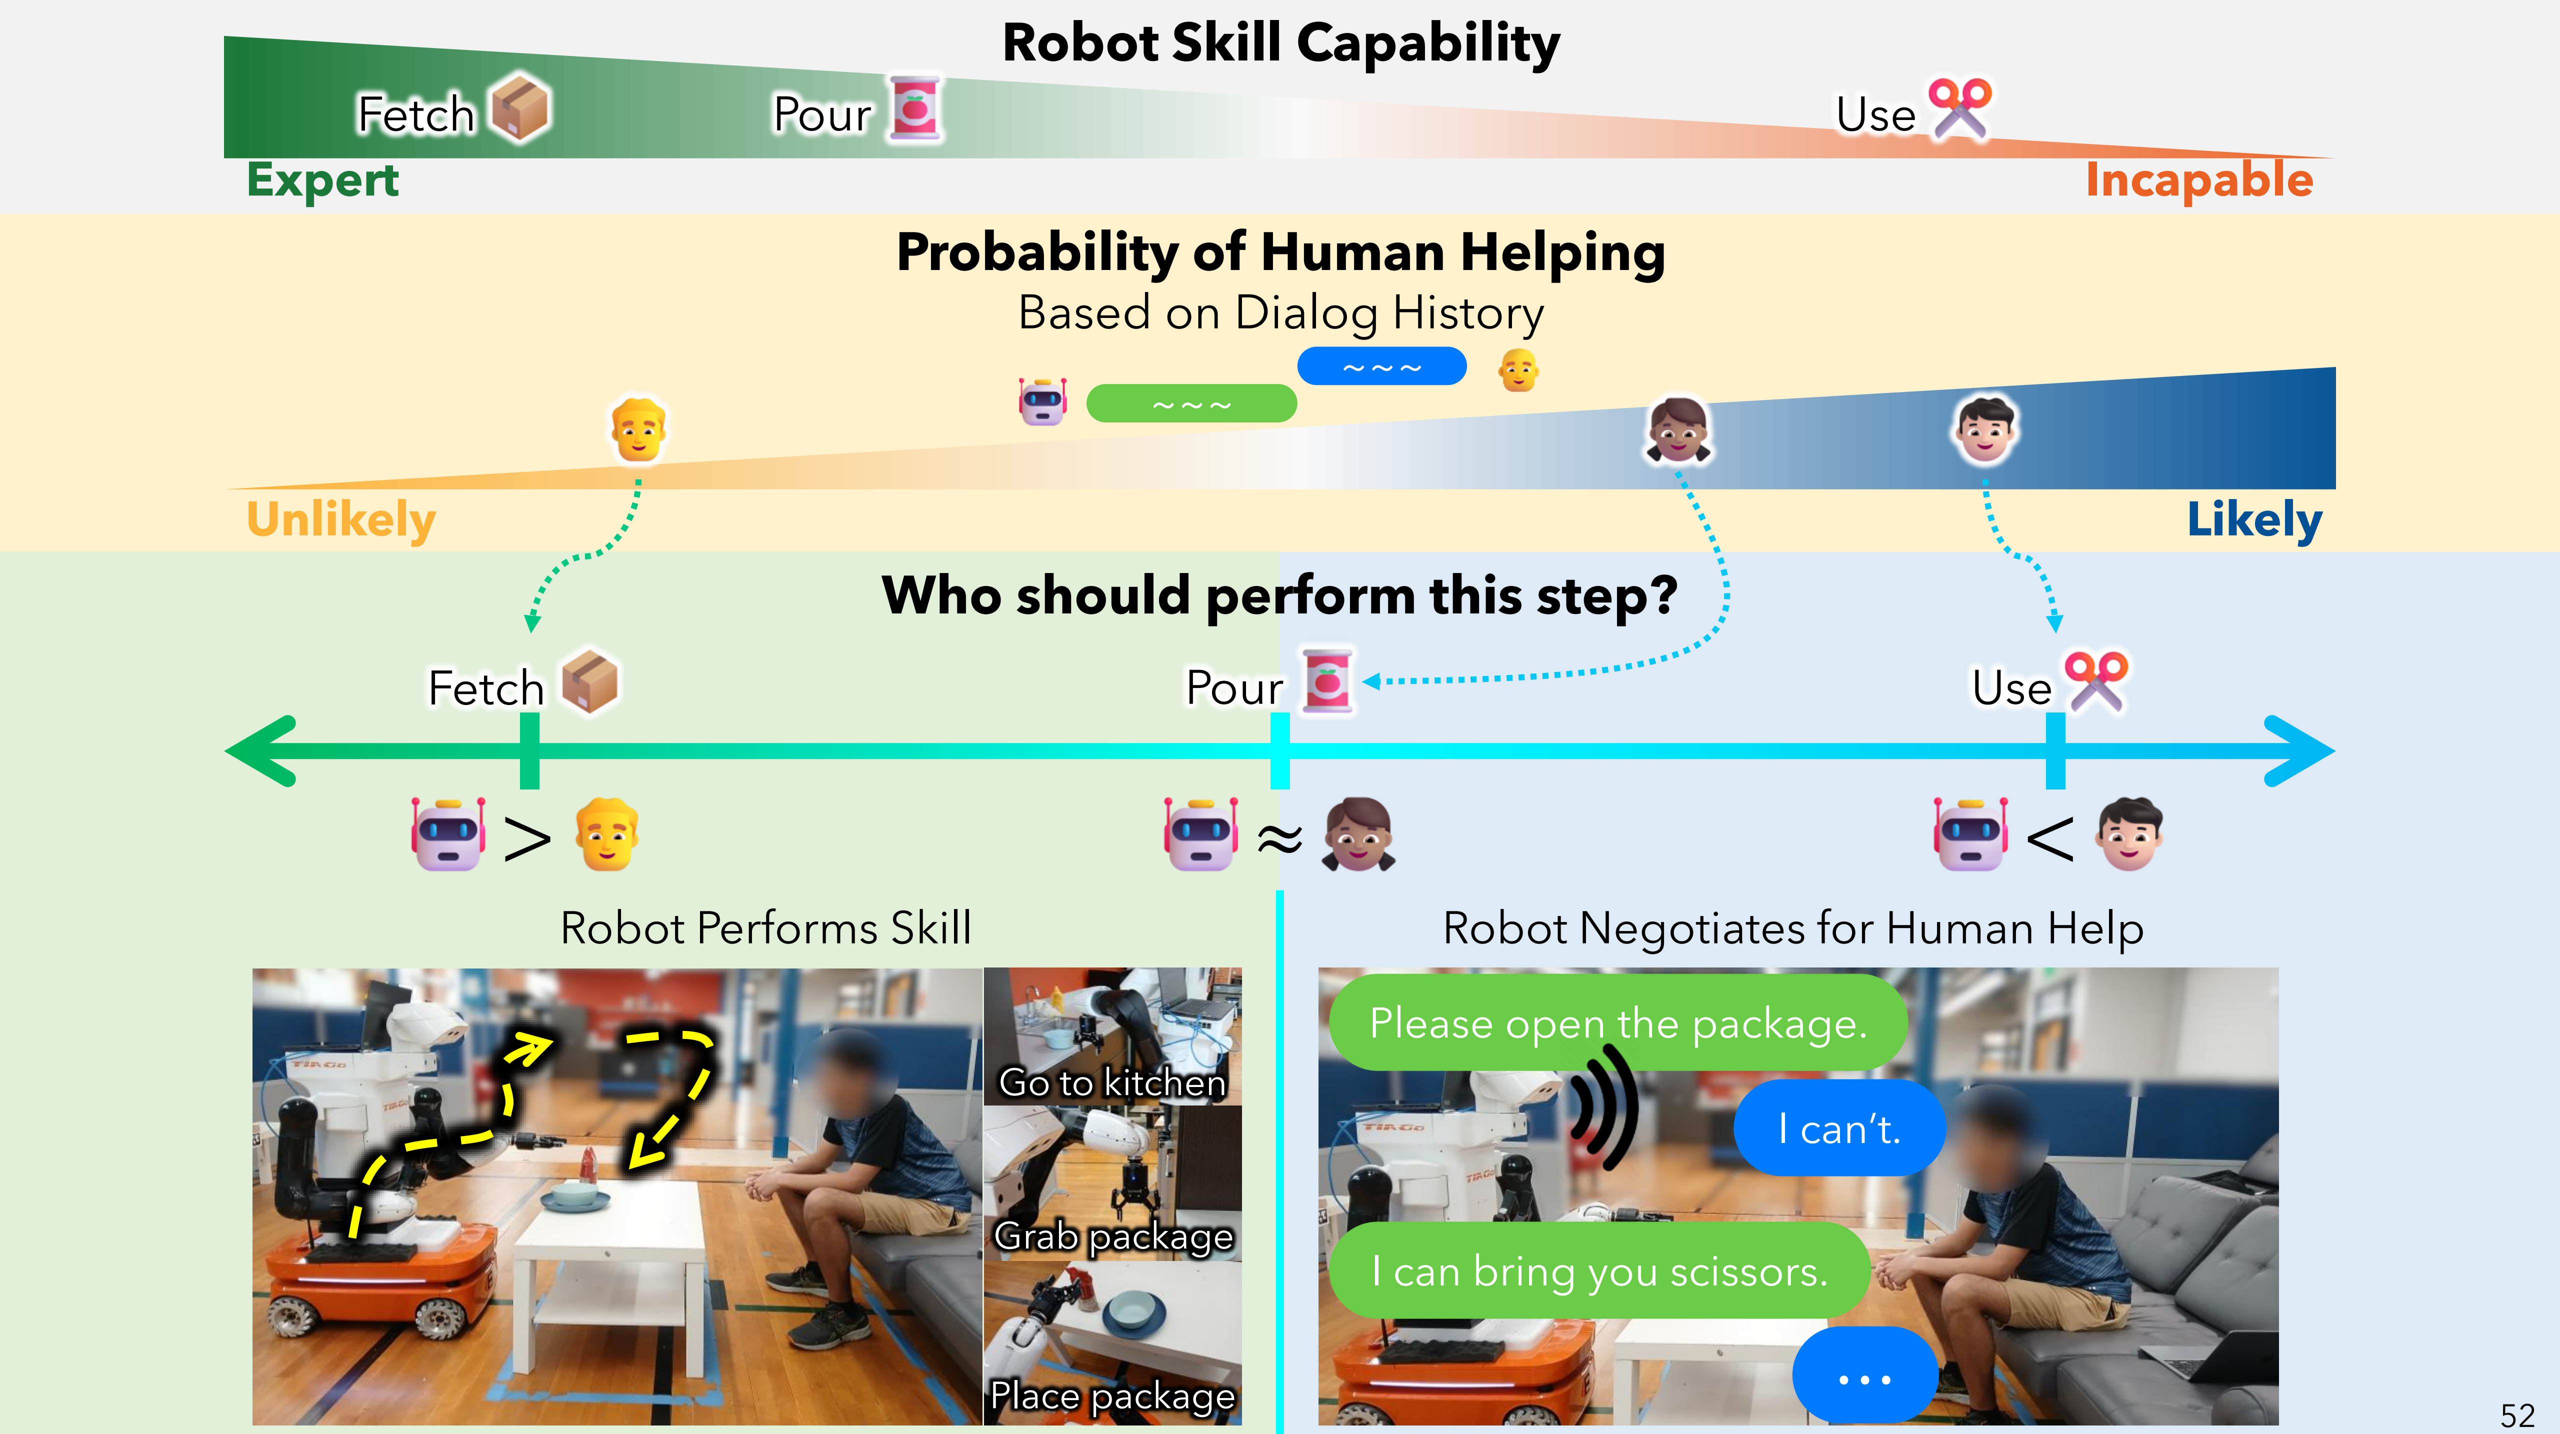
Robot Skill Capability
Fetch
Pour
Use
Expert
Incapable
Probability of Human Helping
Based on Dialog History
~~~
~~~
Unlikely
Likely
Who should perform this step?
Fetch
Pour
Use
>
≈
<
Robot Performs Skill
Go to kitchen
Grab package
Place package
Robot Negotiates for Human Help
Please open the package.
I can't.
I can bring you scissors.
...

Metric: Who does what?

We want $f(a)$ such that:

$f_{\text{🤖}}(a) > f_{\text{👱}}(a) \rightarrow$ 🤖 Performs action a

$f_{\text{🤖}}(a) < f_{\text{👱}}(a) \rightarrow$ 👱 Performs action a

If $\forall t, r_t = -1$, and $\gamma = 1$:

$$Q(s, a) = -\mathbb{E}[\#\text{ timesteps to success on } a \text{ from } s]$$

Inspiration: Temporal Distance RL

Vivek Myers, Chongyi Zheng, Anca Dragan, Sergey Levine, Benjamin Eysenbach. "Learning Temporal Distances: Contrastive Successor Features Can Provide a Metric Structure for Decision-Making."

What else does MicoBot need to decide?

High-Level Strategy

Continue doing task?
Initiate Dialog?
Respond to Dialog?

Propose a way to split the task?

Negotiate for human help?

Task Allocation

Who is best suited to perform each step?

Low-level Actions

What low-level physical actions to execute?

What words to say to the human?

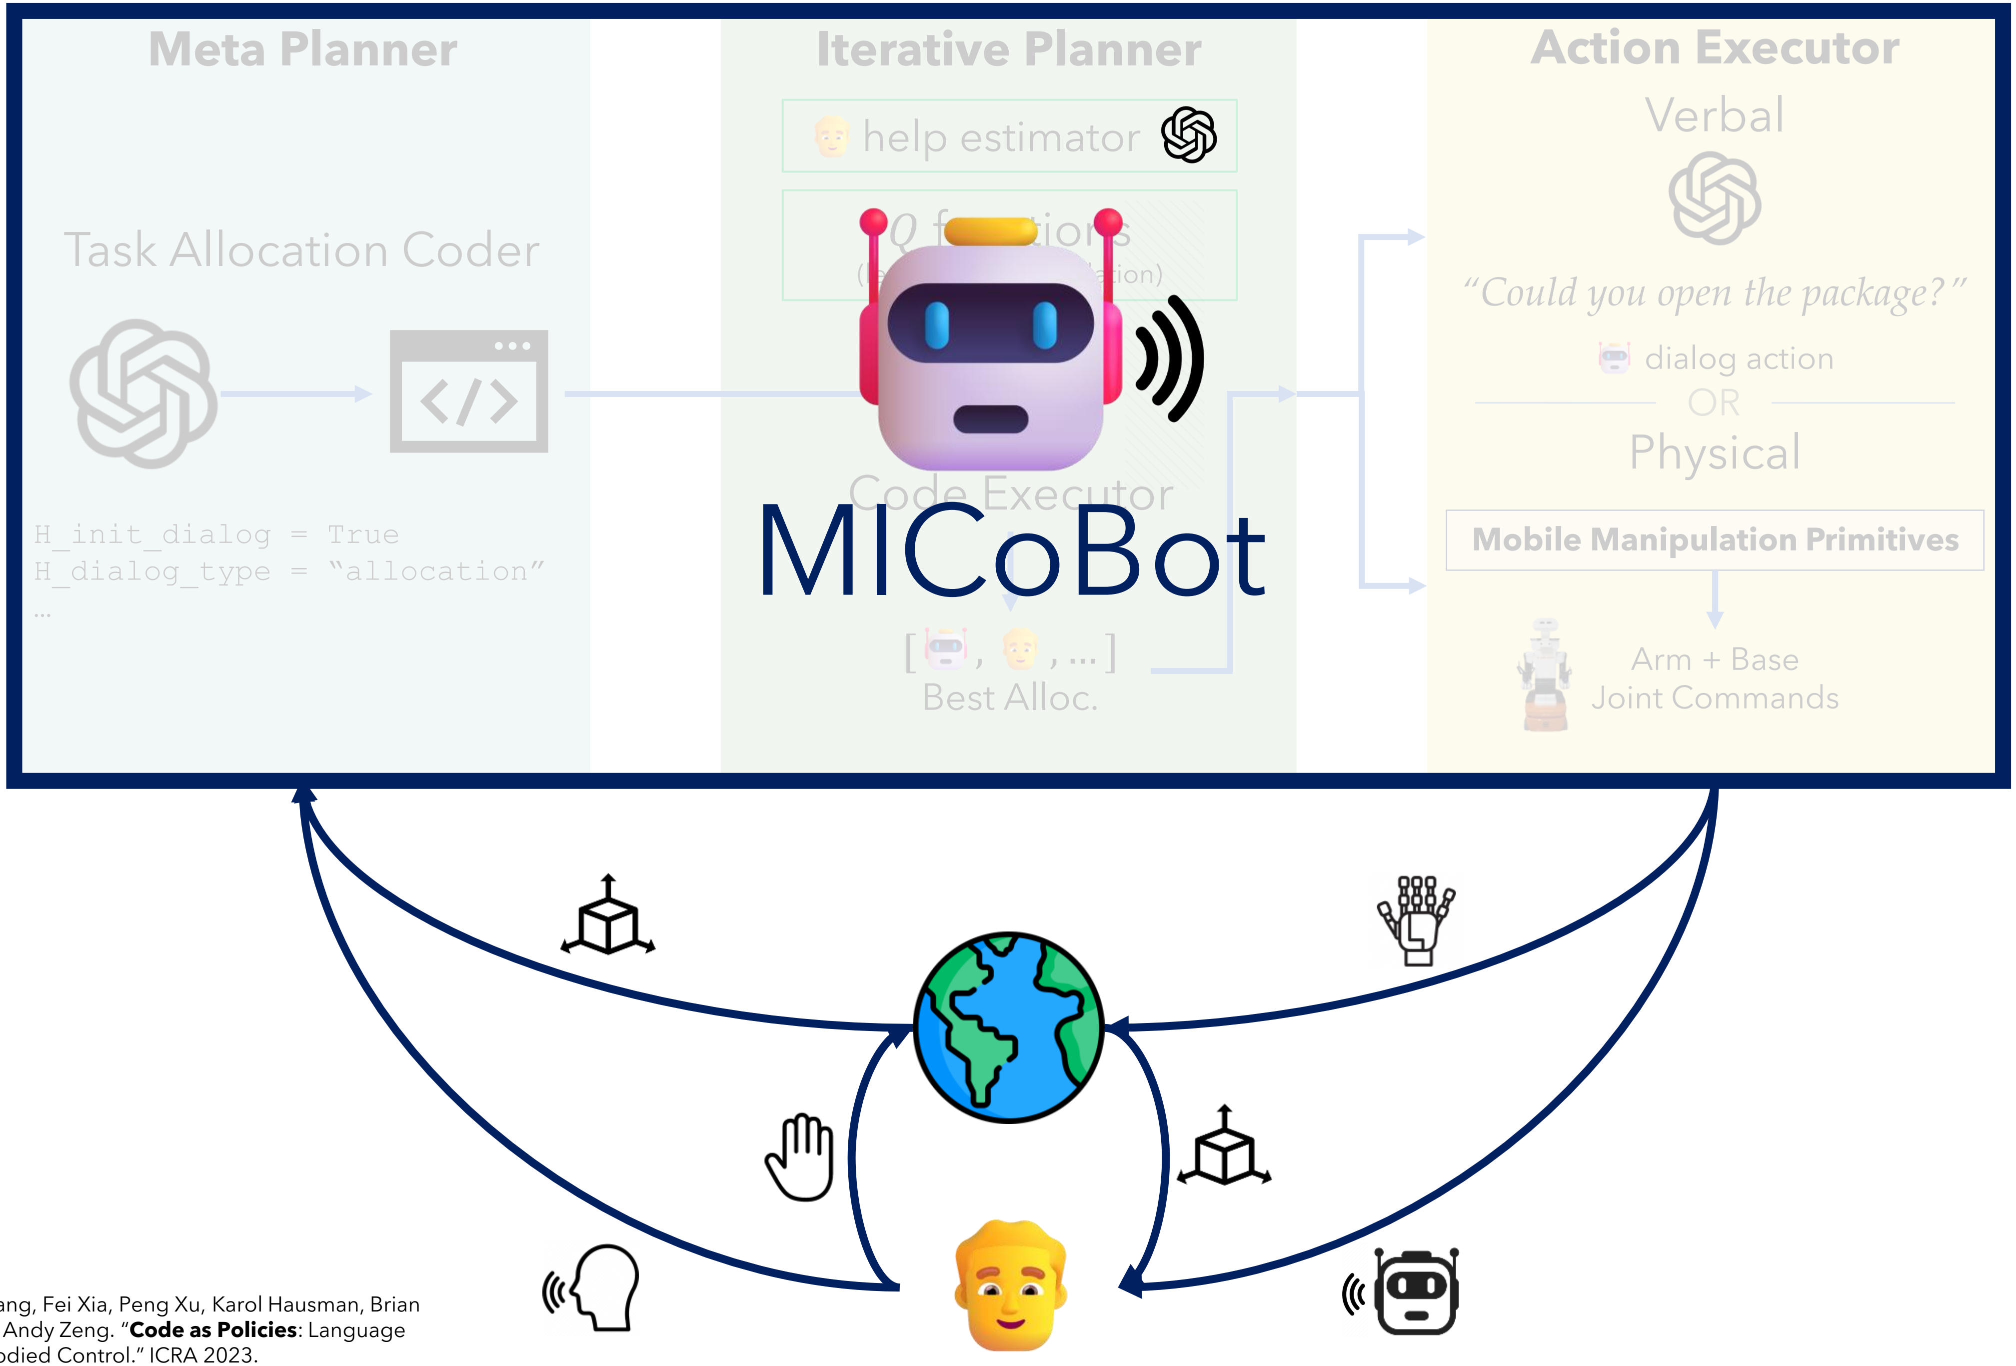

Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. "**Code as Policies**: Language Model Programs for Embodied Control." ICRA 2023.


TIAGo
20x

Real World User Studies (n = 18, 3 tasks)

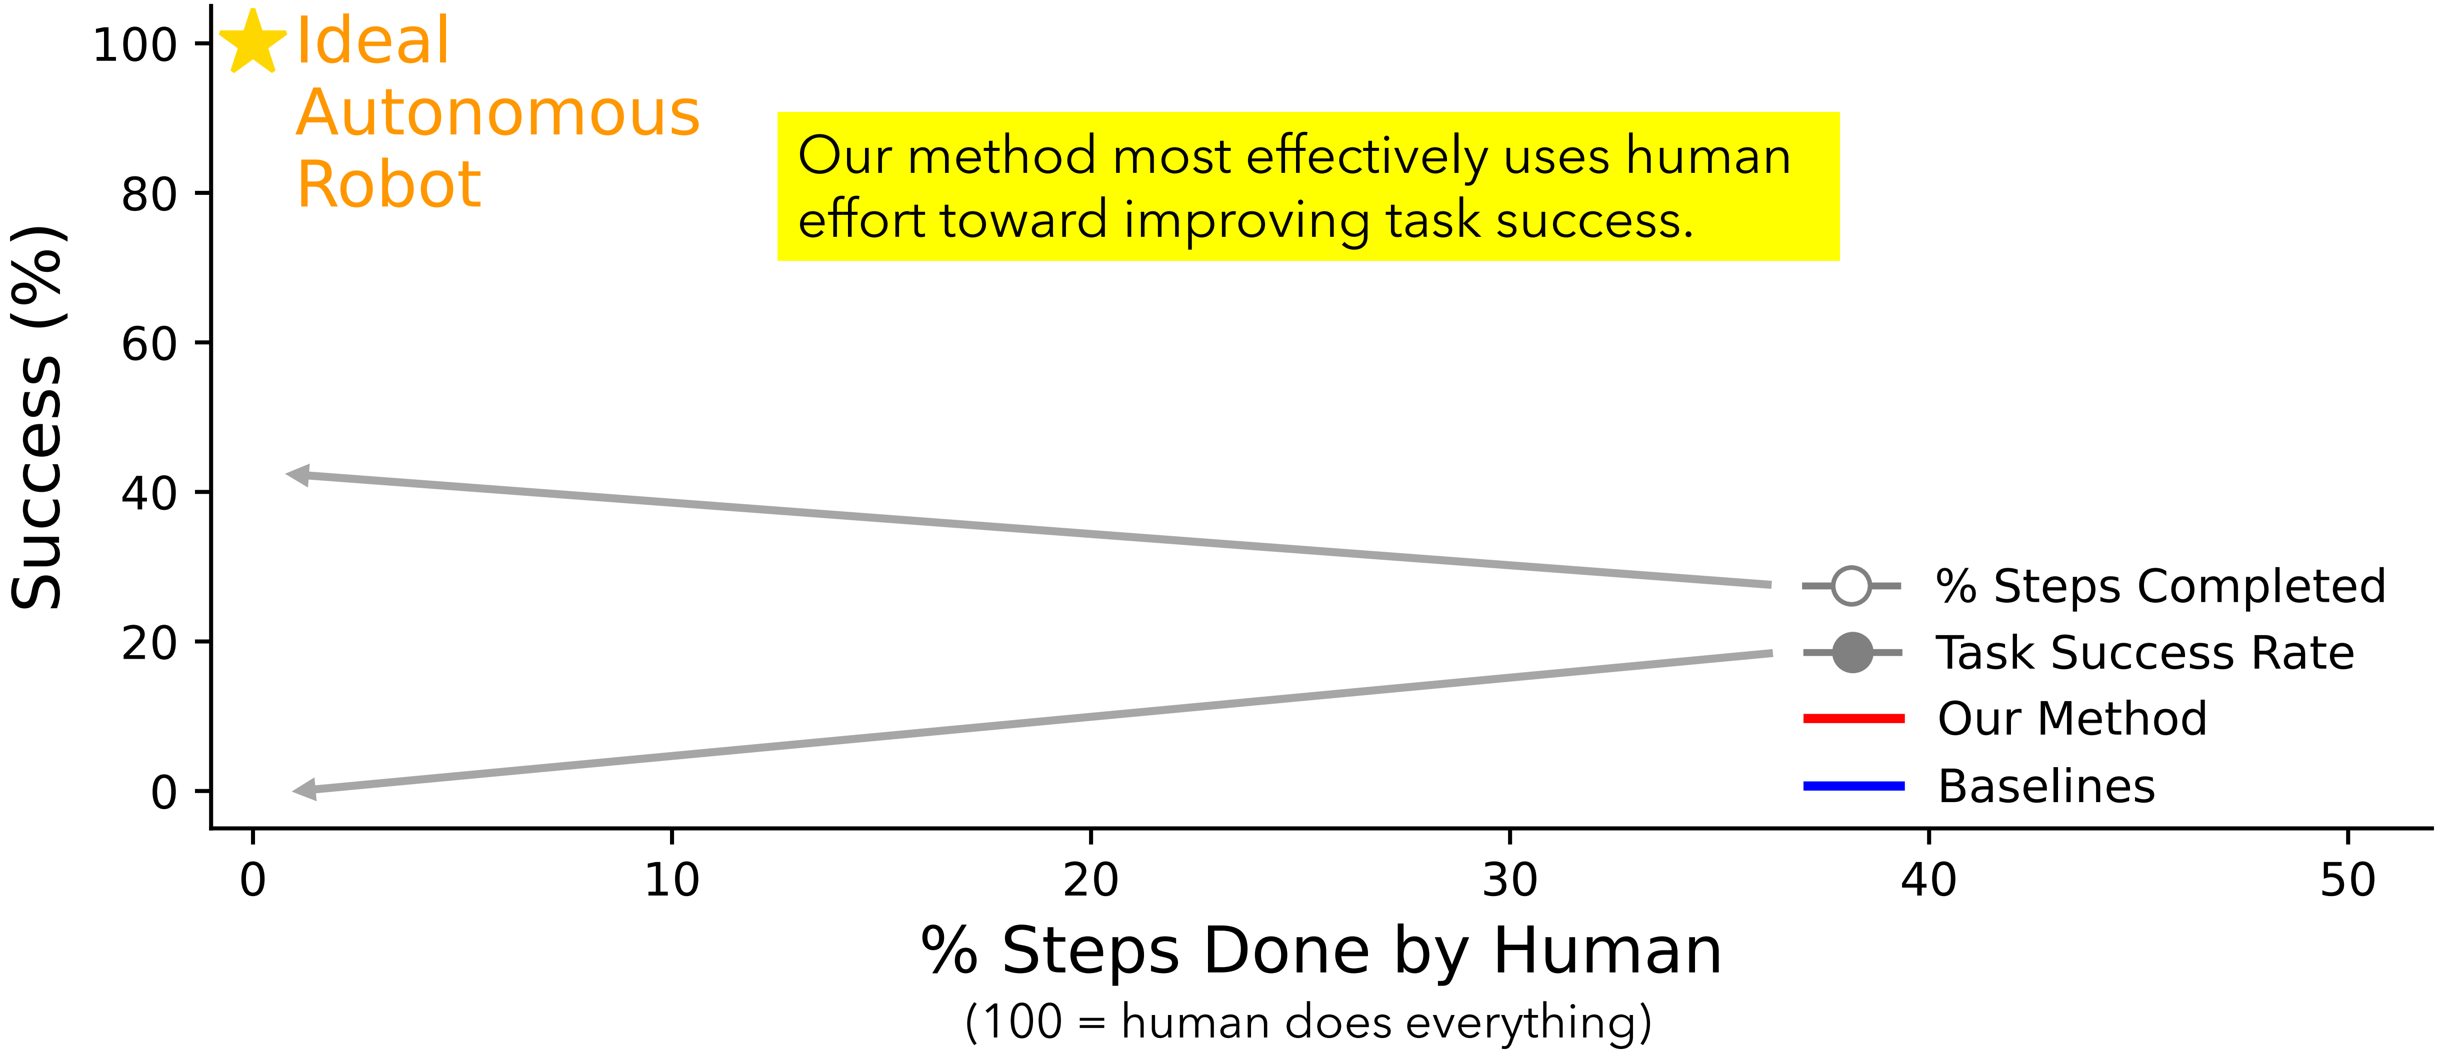

Key Takeaways

1. MicoBot is the first work to enable mixed-initiative dialog for real-world HRI.
2. Mixed-initiative dialog improves collaborative manipulation and adaptiveness compared to single-initiative dialog.

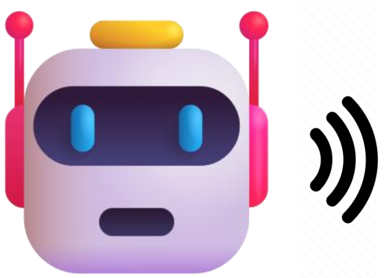

MICoBot

Mixed-initiative dialog
for real-world human+robot manipulation

Outline

Language Improves Robotic Capabilities through…

Semantics for Generalization

1. DeL-TaCo: Simultaneous Learning from Demo + Language
2. Lang4Sim2Real: Language to Bridge the Sim2Real Gap

Communicative Medium

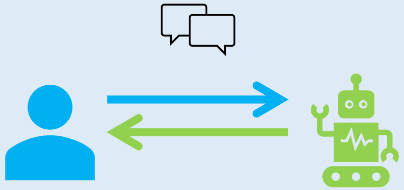

3. MicoBot: Bidirectional Dialog for Human-Robot Collaboration
4. **BookWorm**: Learning to operate new devices with guidebooks and human feedback

Future Work and Conclusion

Transfer Learning in Guidebook Grounding with Human Feedback for Robot Manipulation

Collaborators: Ray Mooney, Roberto Martín-Martín, Will Reger, Cosmo Wu, Da Liu

Under Submission, CoRL 2026

Say you are a housekeeper working in a new home.

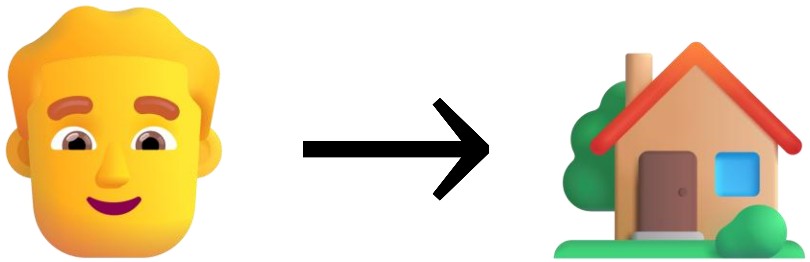

How would you figure out how to perform tasks with new appliances?

Cooking



Making Drinks


Cleaning




Guidebooks
BLACK+DECKER
4-SLICE TOASTER OVEN
Product may vary slightly from what is illustrated.
1. Easy-view glass door
2. Power indicator light
3. Temperature selector knob
4. Toast shade/timer knob with stay on setting
5. Crumb tray
6. Bake pan/drip tray
7. Slide rack/broil rack
8. Extra deep curved interior
HORNO TOSTADOR DE 4 REBANADAS
GRILLE-PAIN FOUR À 4 TRANCHES
amazon basics
Waffle Maker
Gaufrier
Gofrera
B0DM57DNLQ
Cuisinart
Compact Automatic Bread Maker
CBK-110 SERIES
Honeywell Home
X2P Programmable Thermostat
Installation and User Guide
This document contains the following sections:
Package includes:
Mr. Coffee
SIMPLY great COFFEE
PC05 SERIES
USER MANUAL
Hamilton Beach
Breakfast Sandwich Maker
Grille-sandwich pour le déjeuner
Máquina para Preparar Sándwiches de Desayuno
Recipes Included!
CHEFMAN
ROTATING WAFFLE MAKER
USER GUIDE
READ ALL INSTRUCTIONS BEFORE USE

Challenge: Grounding

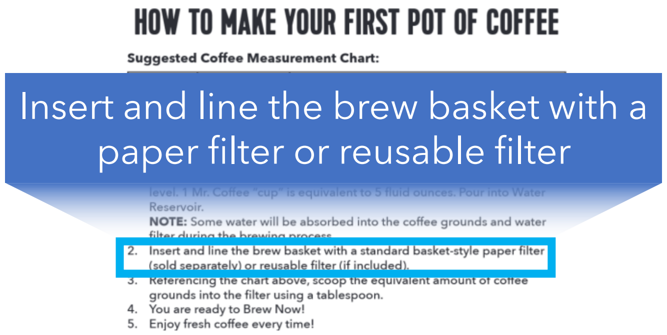

Step-by-Step instructions

Arm movements

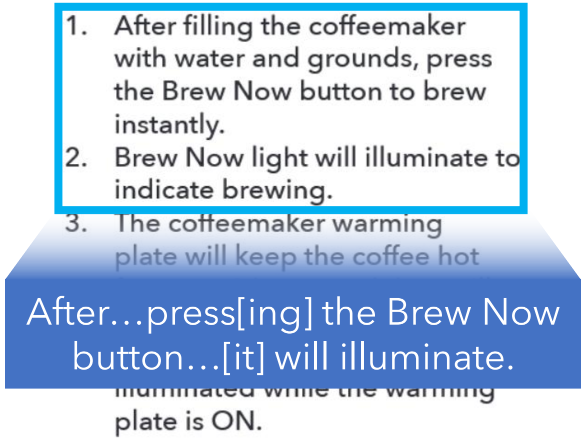

Expected Dynamics

Progress Tracking

Entity 3D Coordinates

Brew Basket: (0.3, 0.2, 0.1)

Hour button: …

Solution: Multi-modal Human Feedback

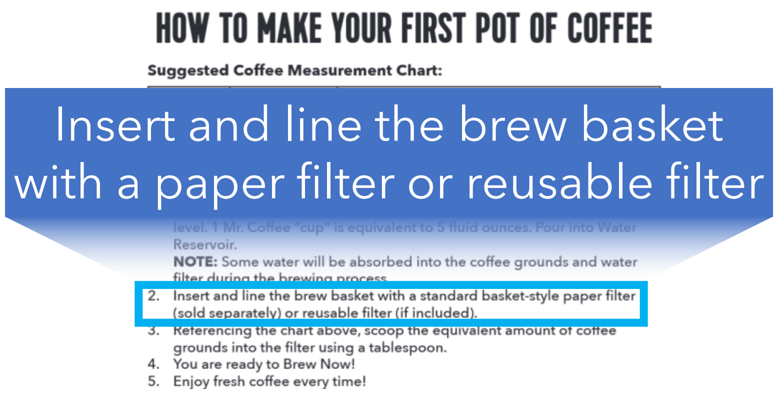
HOW TO MAKE YOUR FIRST POT OF COFFEE

Suggested Coffee Measurement Chart:

level. 1 Mr. Coffee "cup" is equivalent to 5 fluid ounces. Pour into Water Reservoir.

NOTE: Some water will be absorbed into the coffee grounds and water filter during the brewing process.

2. Insert and line the brew basket with a standard basket-style paper filter (sold separately) or reusable filter (if included).
3. Referencing the chart above, scoop the equivalent amount of coffee grounds into the filter using a tablespoon.
4. You are ready to Brew Now!
5. Enjoy fresh coffee every time!

Insert and line the brew basket with a paper filter or reusable filter

Why is the filter falling out?

Move your arm lower.

Arm trajectory adjustments


GET TO KNOW YOUR COFFEE MAKER

1. Brew Basket
2. Water Reservoir
3. Brew Later Button
4. Brew Now Button
5. Clock Display
6. Hour/Minute Buttons
7. Carafe
8. Shower Head
9. Warming Plate

UNDERSTAND YOUR CONTROL PANEL

A. Clock Display
B. AM/PM indicator
C. Brew Now Indicator
D. Brew Now Button
E. Minute Button
F. Hour Button
G. Brew Later Button
H. Brew Later Indicator

I can't find the brew basket.

I've circled it.

Entity 3D Coordinates

Brew Basket: (0.3, 0.2, 0.1)

Hour button: …

Human Feedback is Costly.
What else can we leverage?

Experience with previous similar devices

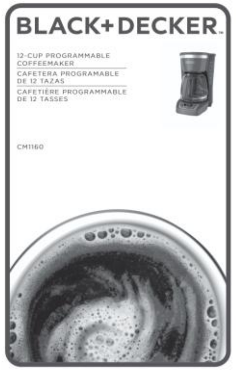

Learn new devices faster by transferring knowledge from previous, similar devices.

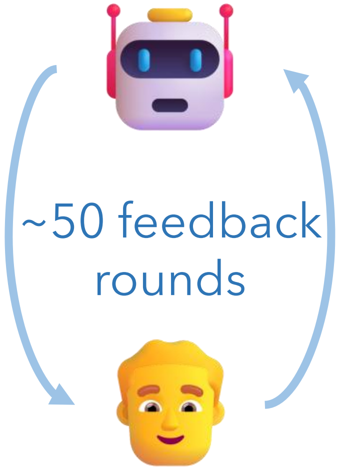

Knowledge Transfer

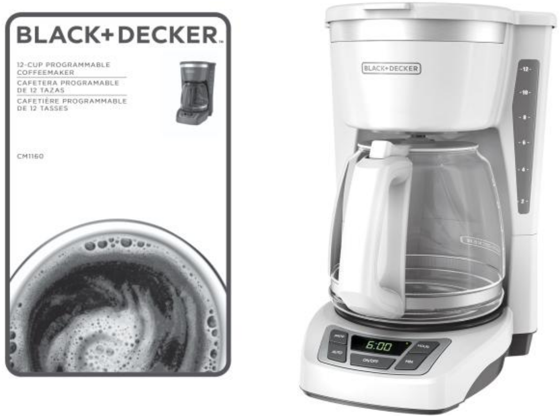

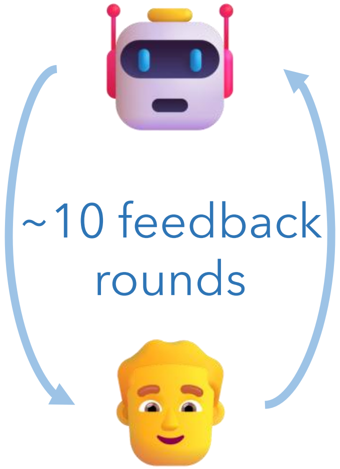

Key insight: Guidebooks provide structural priors for action generation and environmental dynamics.

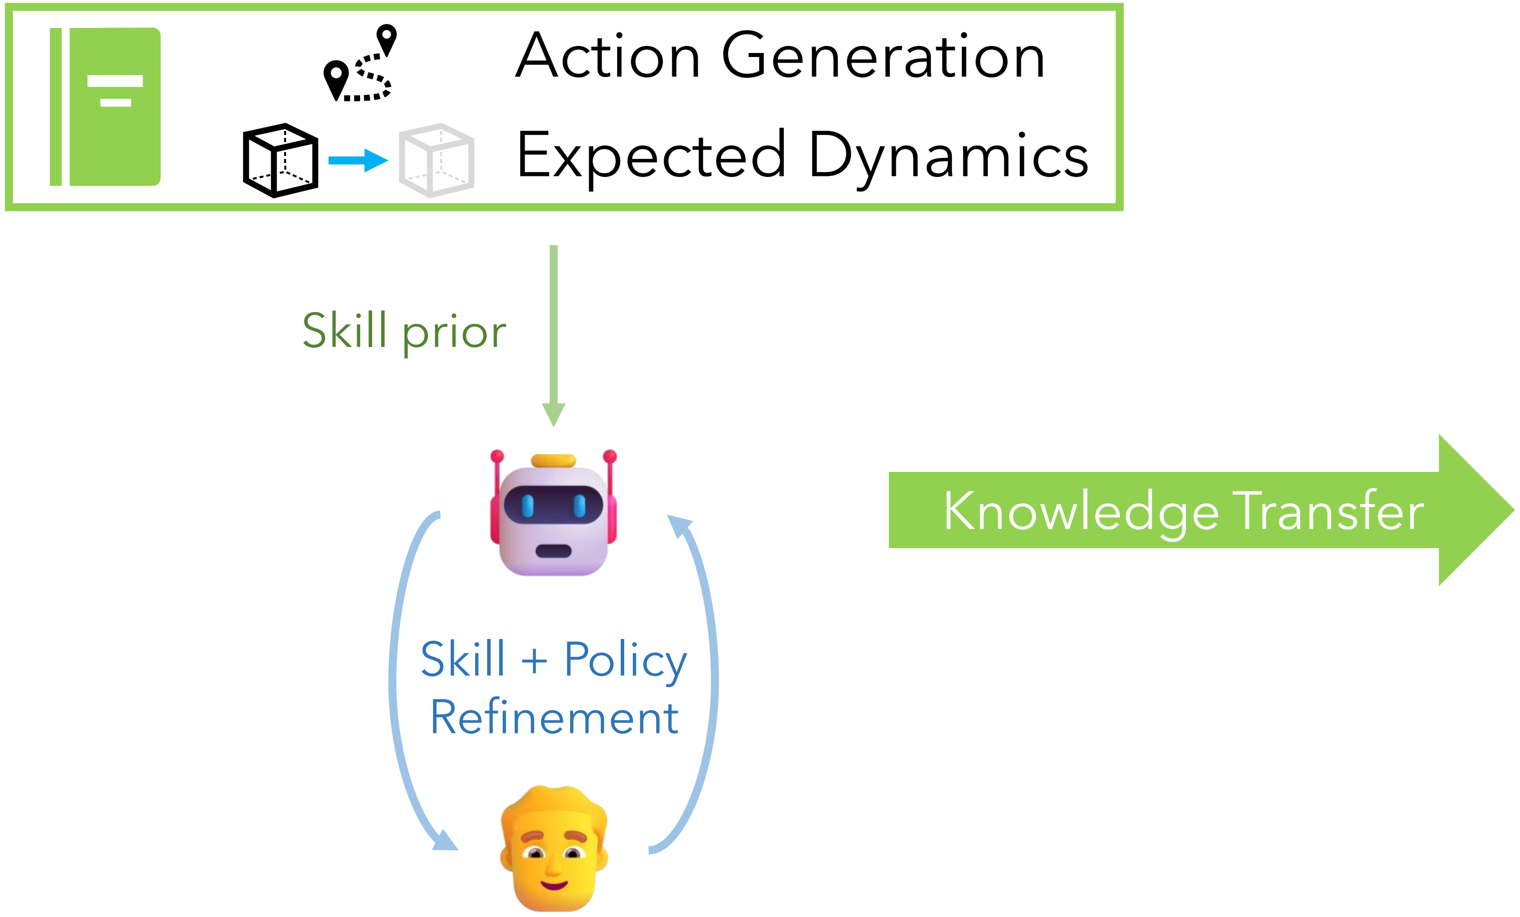

These priors can be refined with human feedback and transferred to efficiently learn from new guidebooks of related devices.

Prior Work: Learning from Manuals

Learning to Play Civilization II from Manuals
(Branavan+, JAIR 2012)

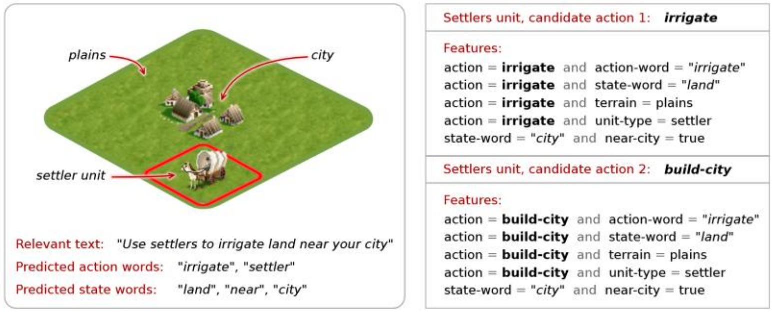

Unable to ground actions and entities to the real world.

IKEA assembly from visual diagrams
(Zhang+, 2025)

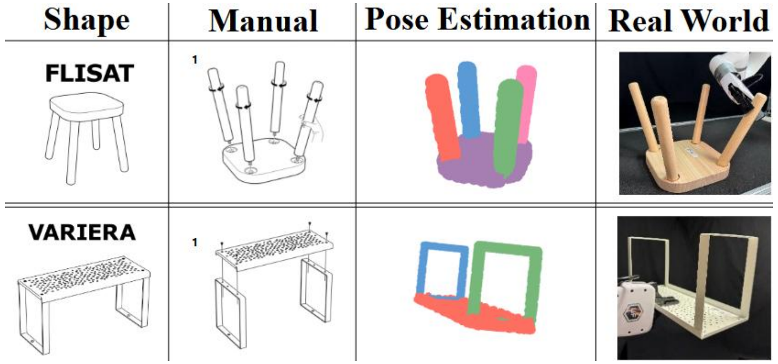

Cannot leverage rich textual info from guidebooks.

Symbolic Dynamics Models from Manuals
(Zhang+, 2025)

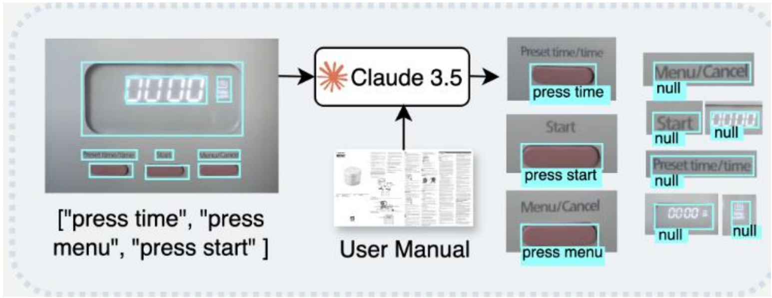

Does not learn or refine skills needed for manipulation.

SRK Branavan, David Silver, Regina Barzilay. "Learning to Win by Reading Manuals in a Monte-Carlo Framework." JAIR 2012.
Jian Zhang, Hanbo Zhang, Anxing Xiao, David Hsu. "Robot Operation of Home Appliances by Reading User Manuals." arxiv 2025.
Chenrui Tie, Shengxiang Sun, Jinxuan Zhu, Yiwei Liu, Jingxiang Guo, Yue Hu, Haonan Chen, Junting Chen, Ruihai Wu, Lin Shao. "Manual2Skill: Learning to Read Manuals and Acquire Robotic Skills for Furniture Assembly Using Vision-Language Models." arxiv 2025.

BookWorm

Bootstrapping **O**bject-grounded **O**perational **K**nowledge
With **O**n-demand **R**obot **M**entoring

Extracting Transferrable Action and Dynamics Priors from Guidebooks with Human Refinement

What does BookWorm need to do?

Acquire task-relevant skills based on the guidebook and task
Determine where device entities are

Predict the correct skill to execute

Ask for human feedback when skills fail
Update behavior with human feedback

How BookWorm Works

Acquire task-relevant skills based on the guidebook and task
Determine where device entities are

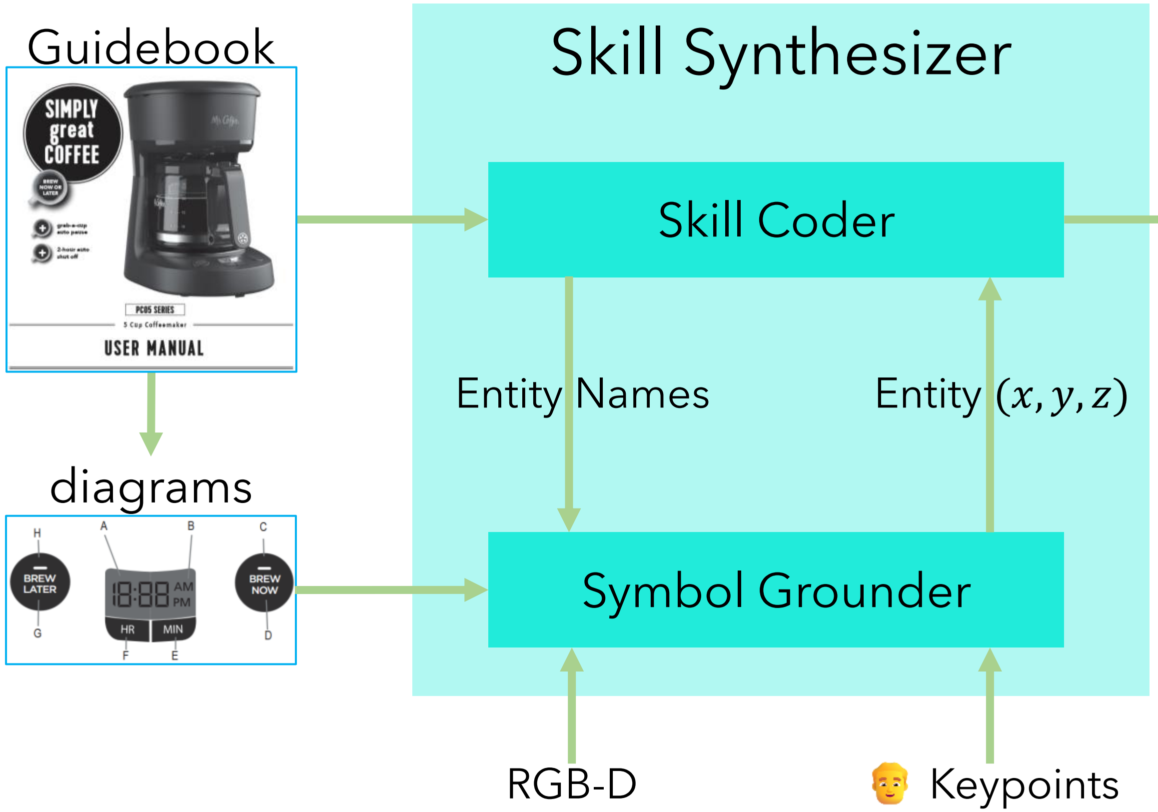

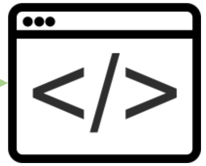

Code-as-Policies-style python skill functions

```
def grasp_and_place(…):
def change_lid_state(…):
…
```

Functions based on bare-bones API:

```
move_to_pose(…)
close_gripper(…)
open_gripper(…)
press_until_force(…)
```

Code-as-Policies: Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, Andy Zeng. "Code as Policies: Language Model Programs for Embodied Control."

Predict the correct skill to execute

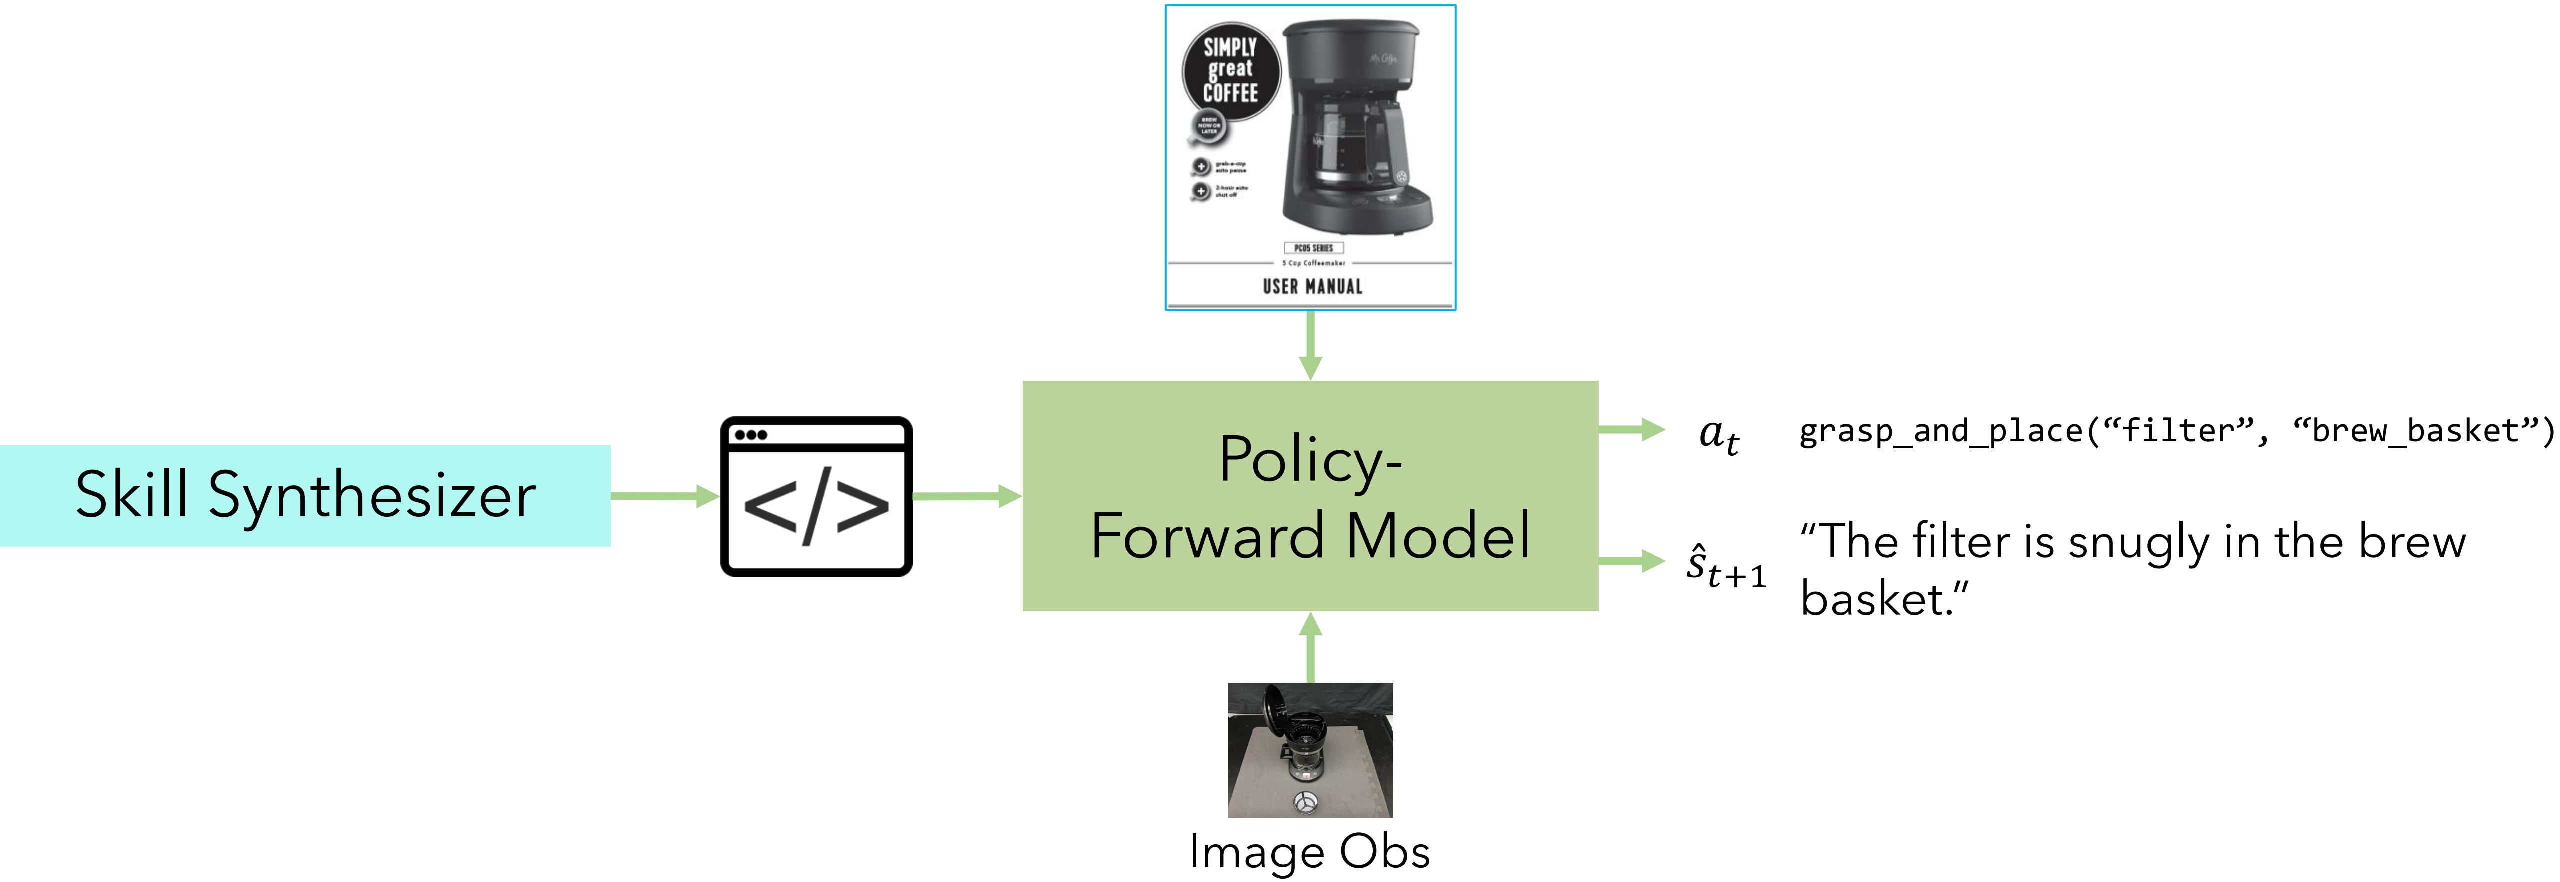

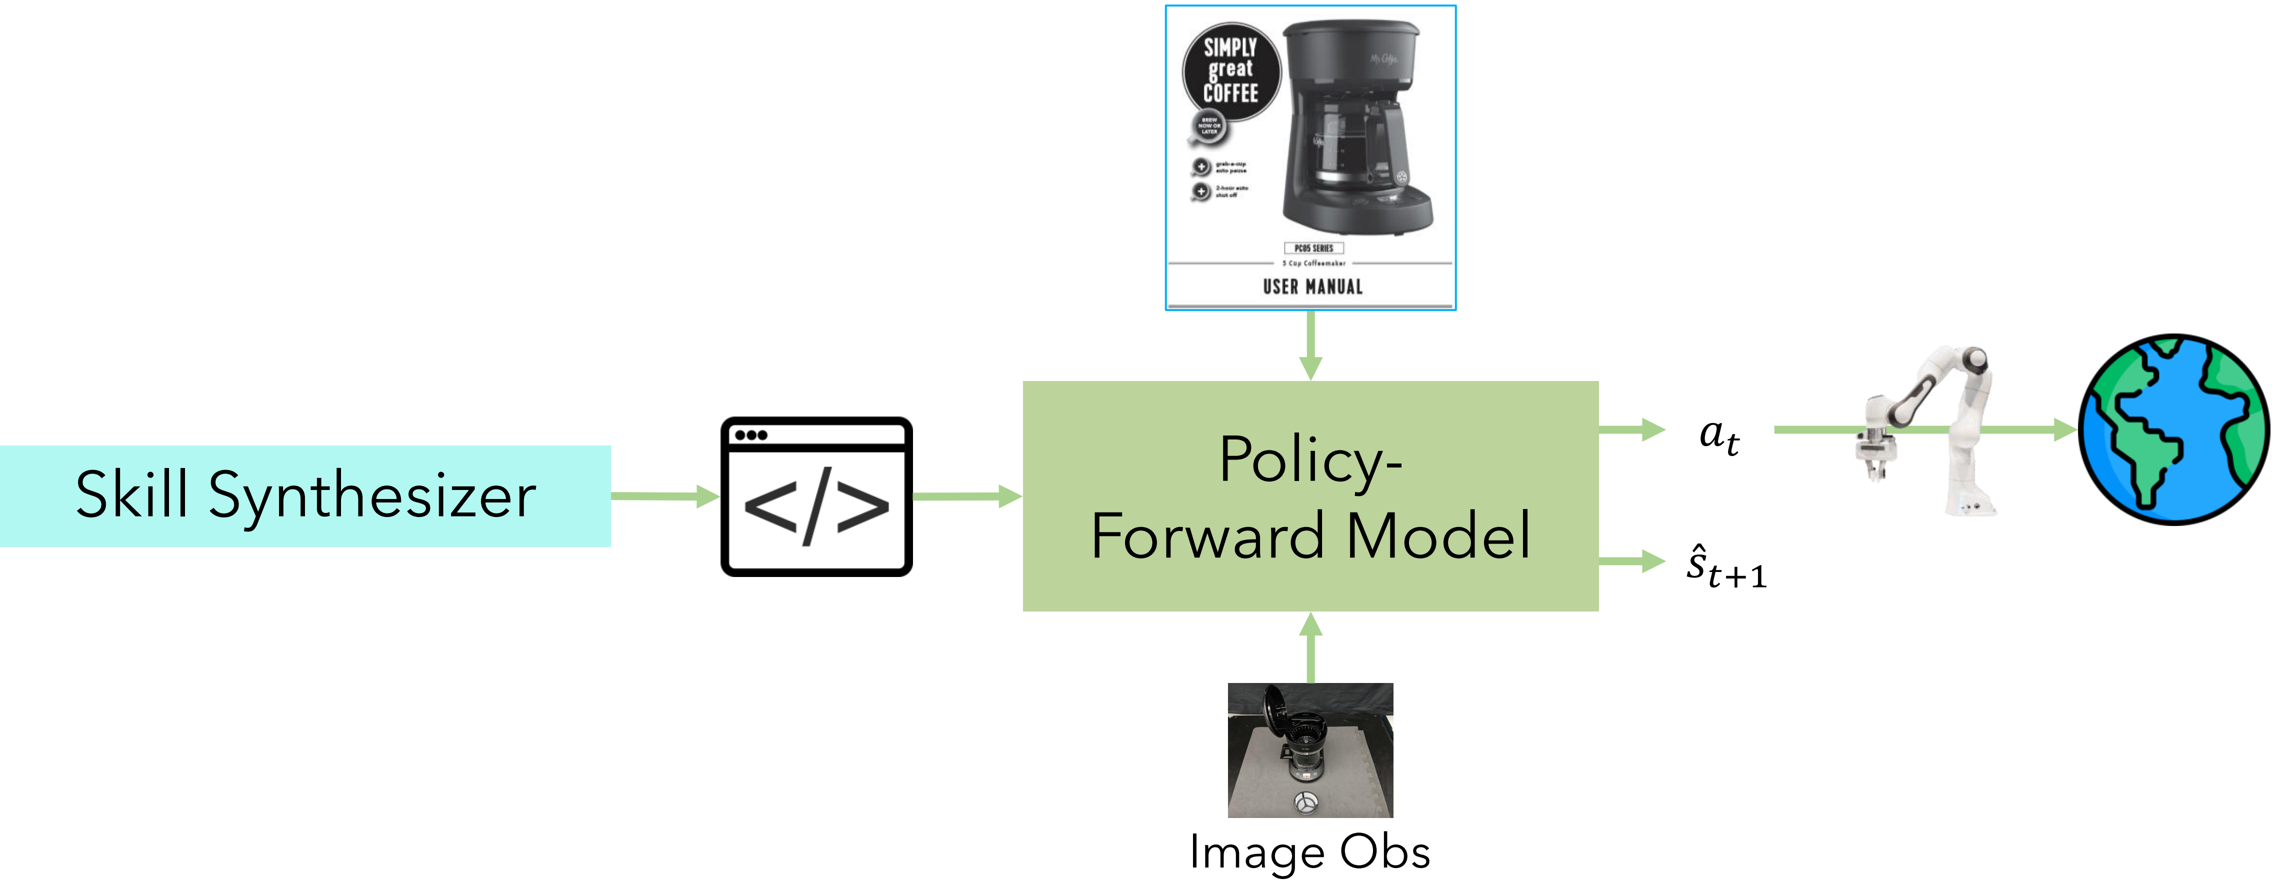
SIMPLY great COFFEE
PC05 SERIES
5 Cup Coffeemaker
USER MANUAL
Skill Synthesizer
Policy-
Forward Model
a_t
$\hat{s}_{t+1}$
Image Obs


10x

Ask for human feedback when skills fail
Update behavior with human feedback

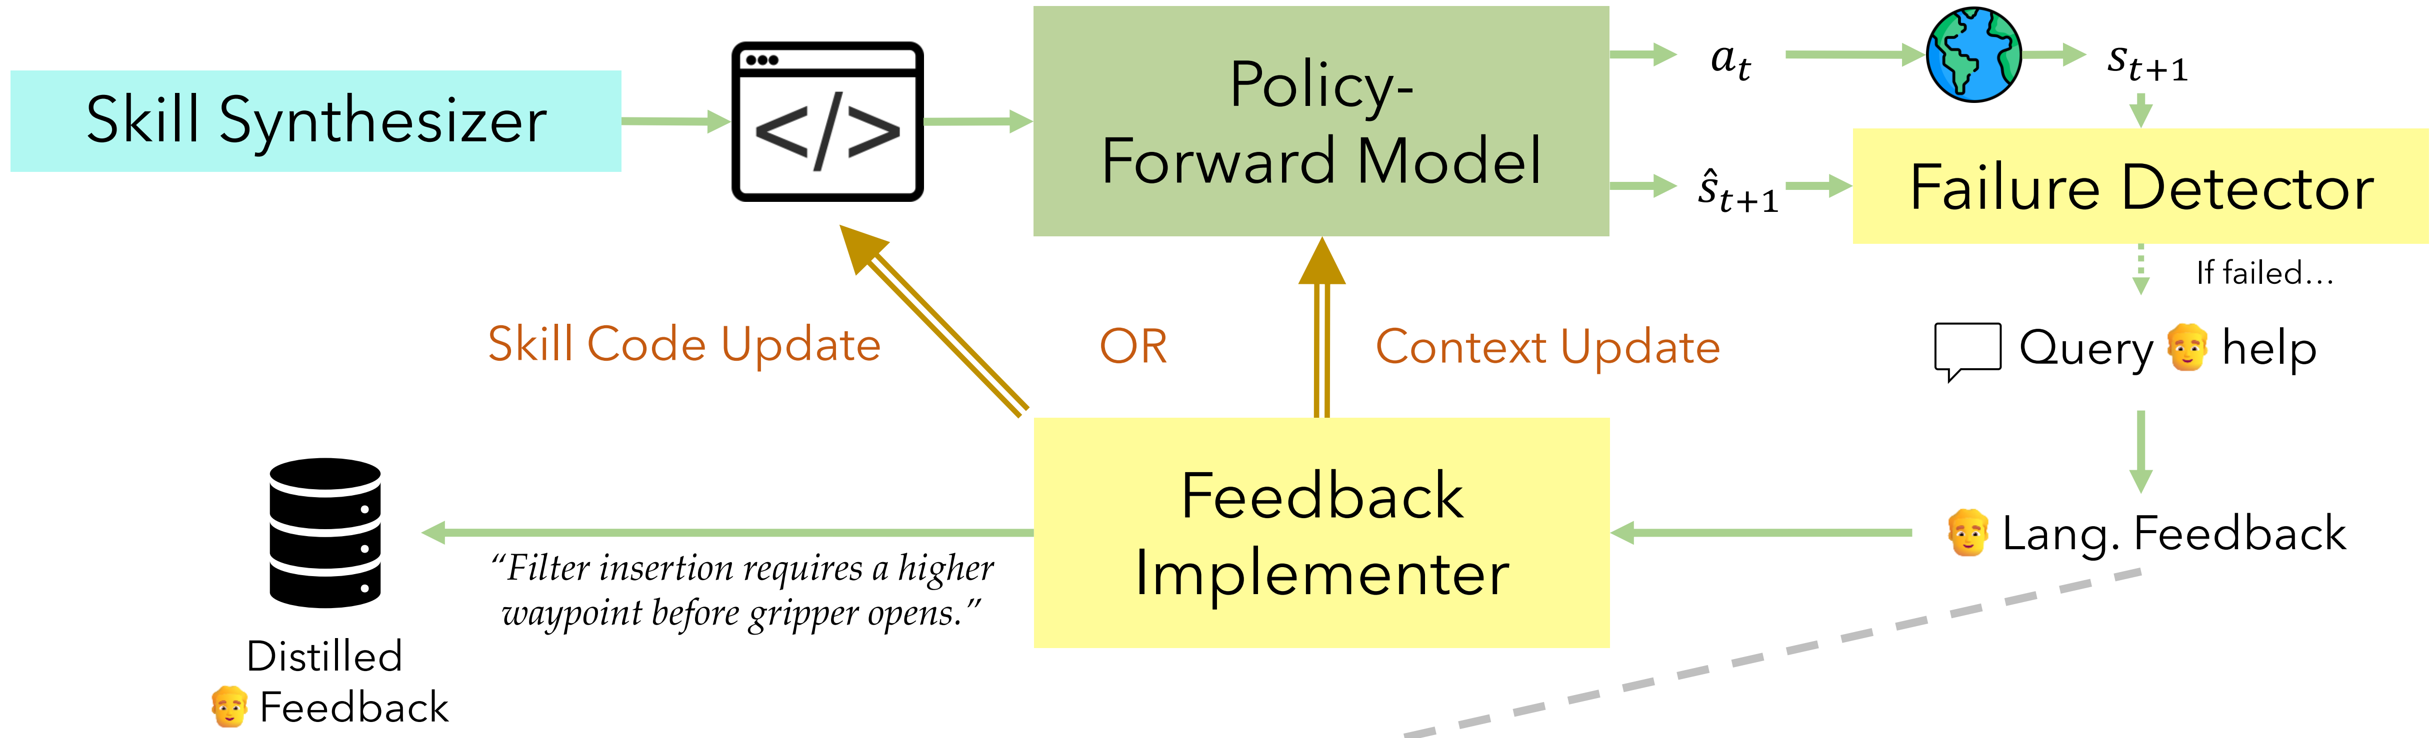

:*"Try to get the filter higher (+z direction) before putting the filter in."*


10x


10x

Experimental Devices

Can BookWorm learn to operate a **single device**?

“Turn off schedule settings.”

“Brew coffee now and adjust time by 2 hours.”

Can BookWorm **transfer** knowledge from one device to learn a new one faster?

“Brew coffee now and adjust time by 2 hours.”

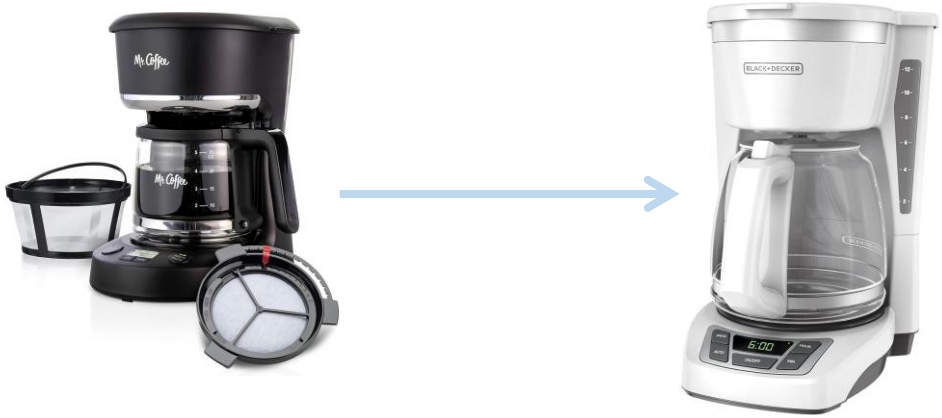

“Make a light brown waffle.”

How does BookWorm compare to prior code-based policy baselines?

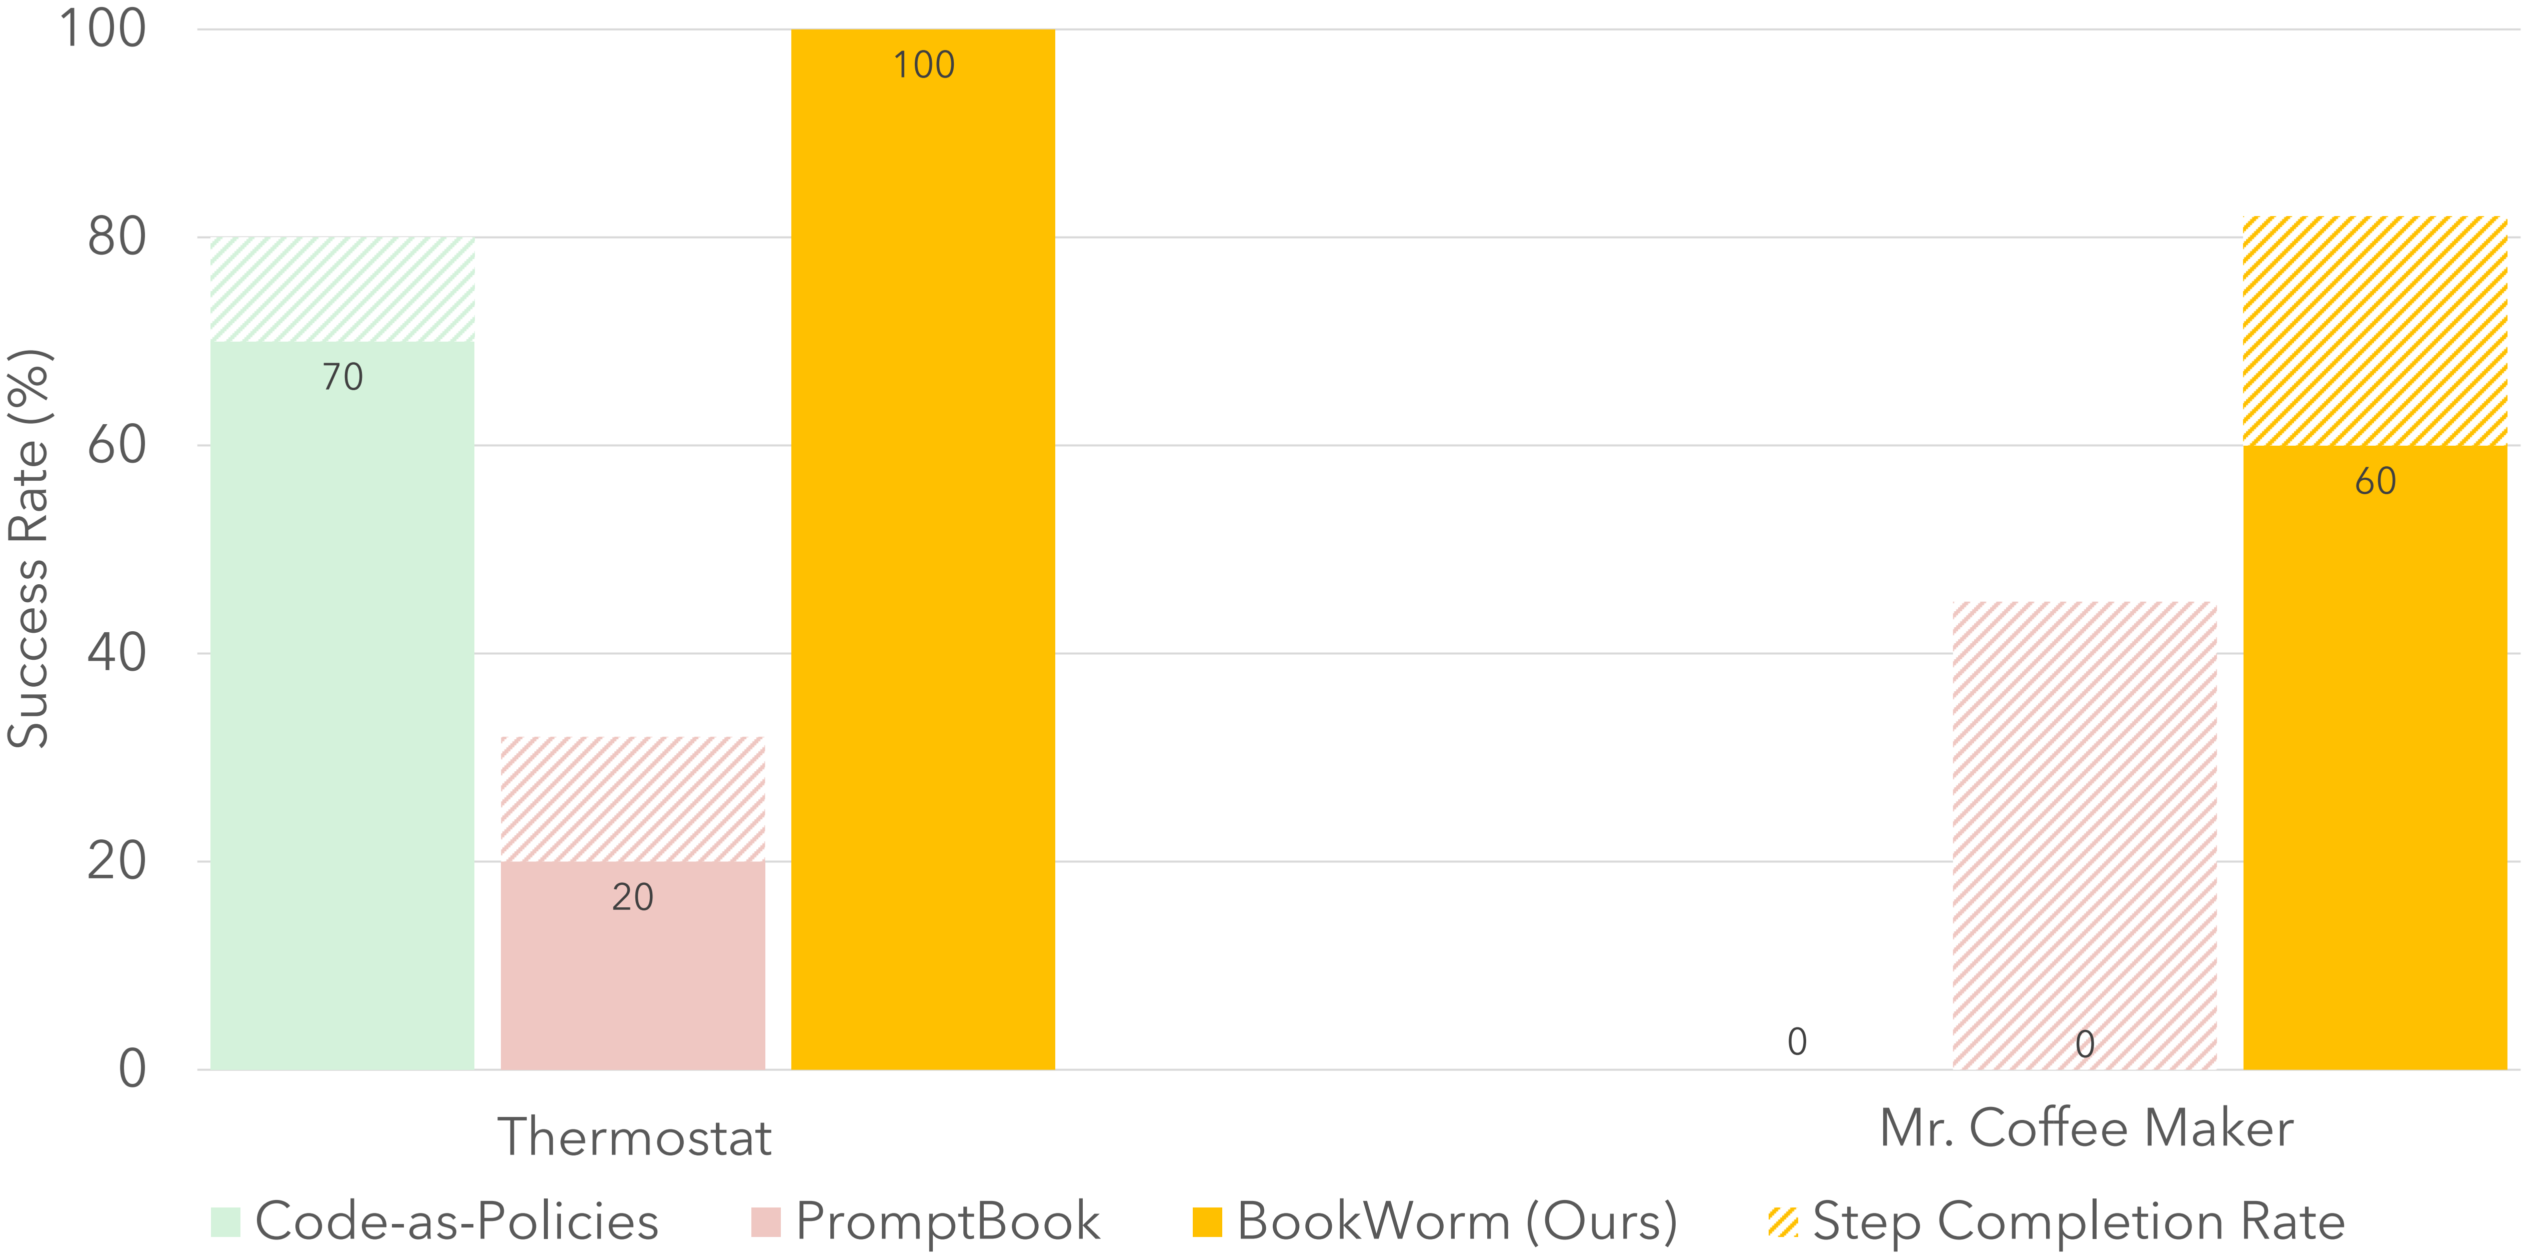

(n=10 / bar)

PromptBook: Montserrat Gonzalez Arenas, Ted Xiao, Sumeet Singh, Vidhi Jain, Allen Z. Ren, Quan Vuong, Jake Varley, Alexander Herzog, Isabel Leal, Sean Kirmani, Dorsa Sadigh, Vikas Sindhwani, Kanishka Rao, Jacky Liang, Andy Zeng“**How to Prompt Your Robot: A PromptBook for Manipulation Skills with Code as Policies.”**

How important are individual components of BookWorm?

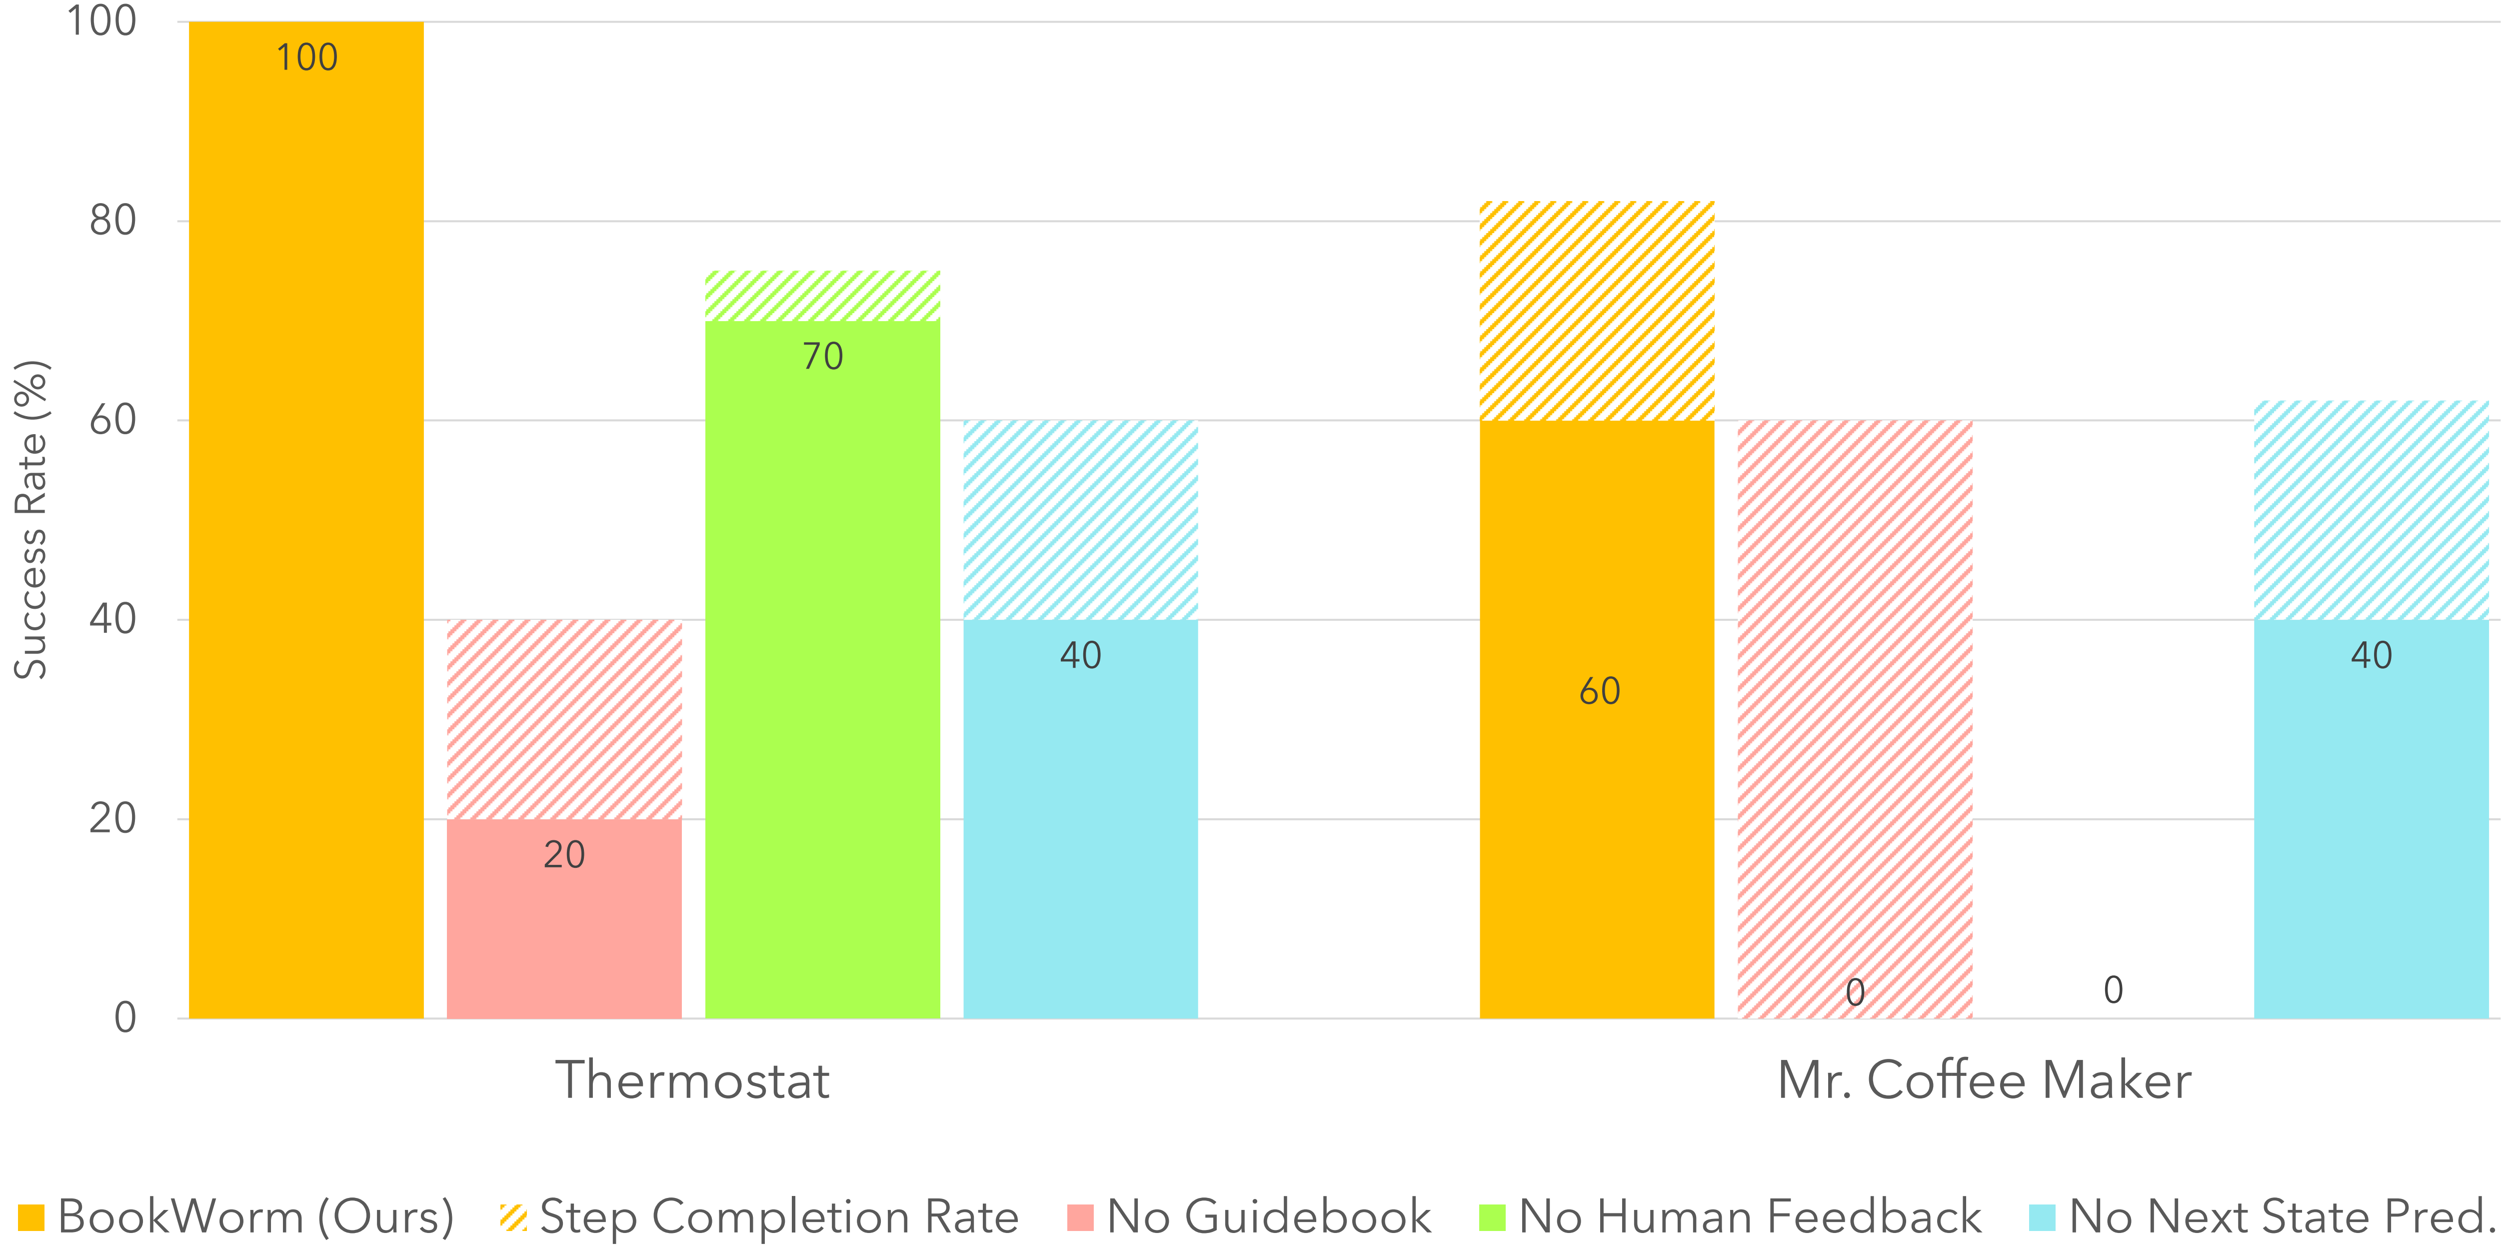

(n=10 / bar)

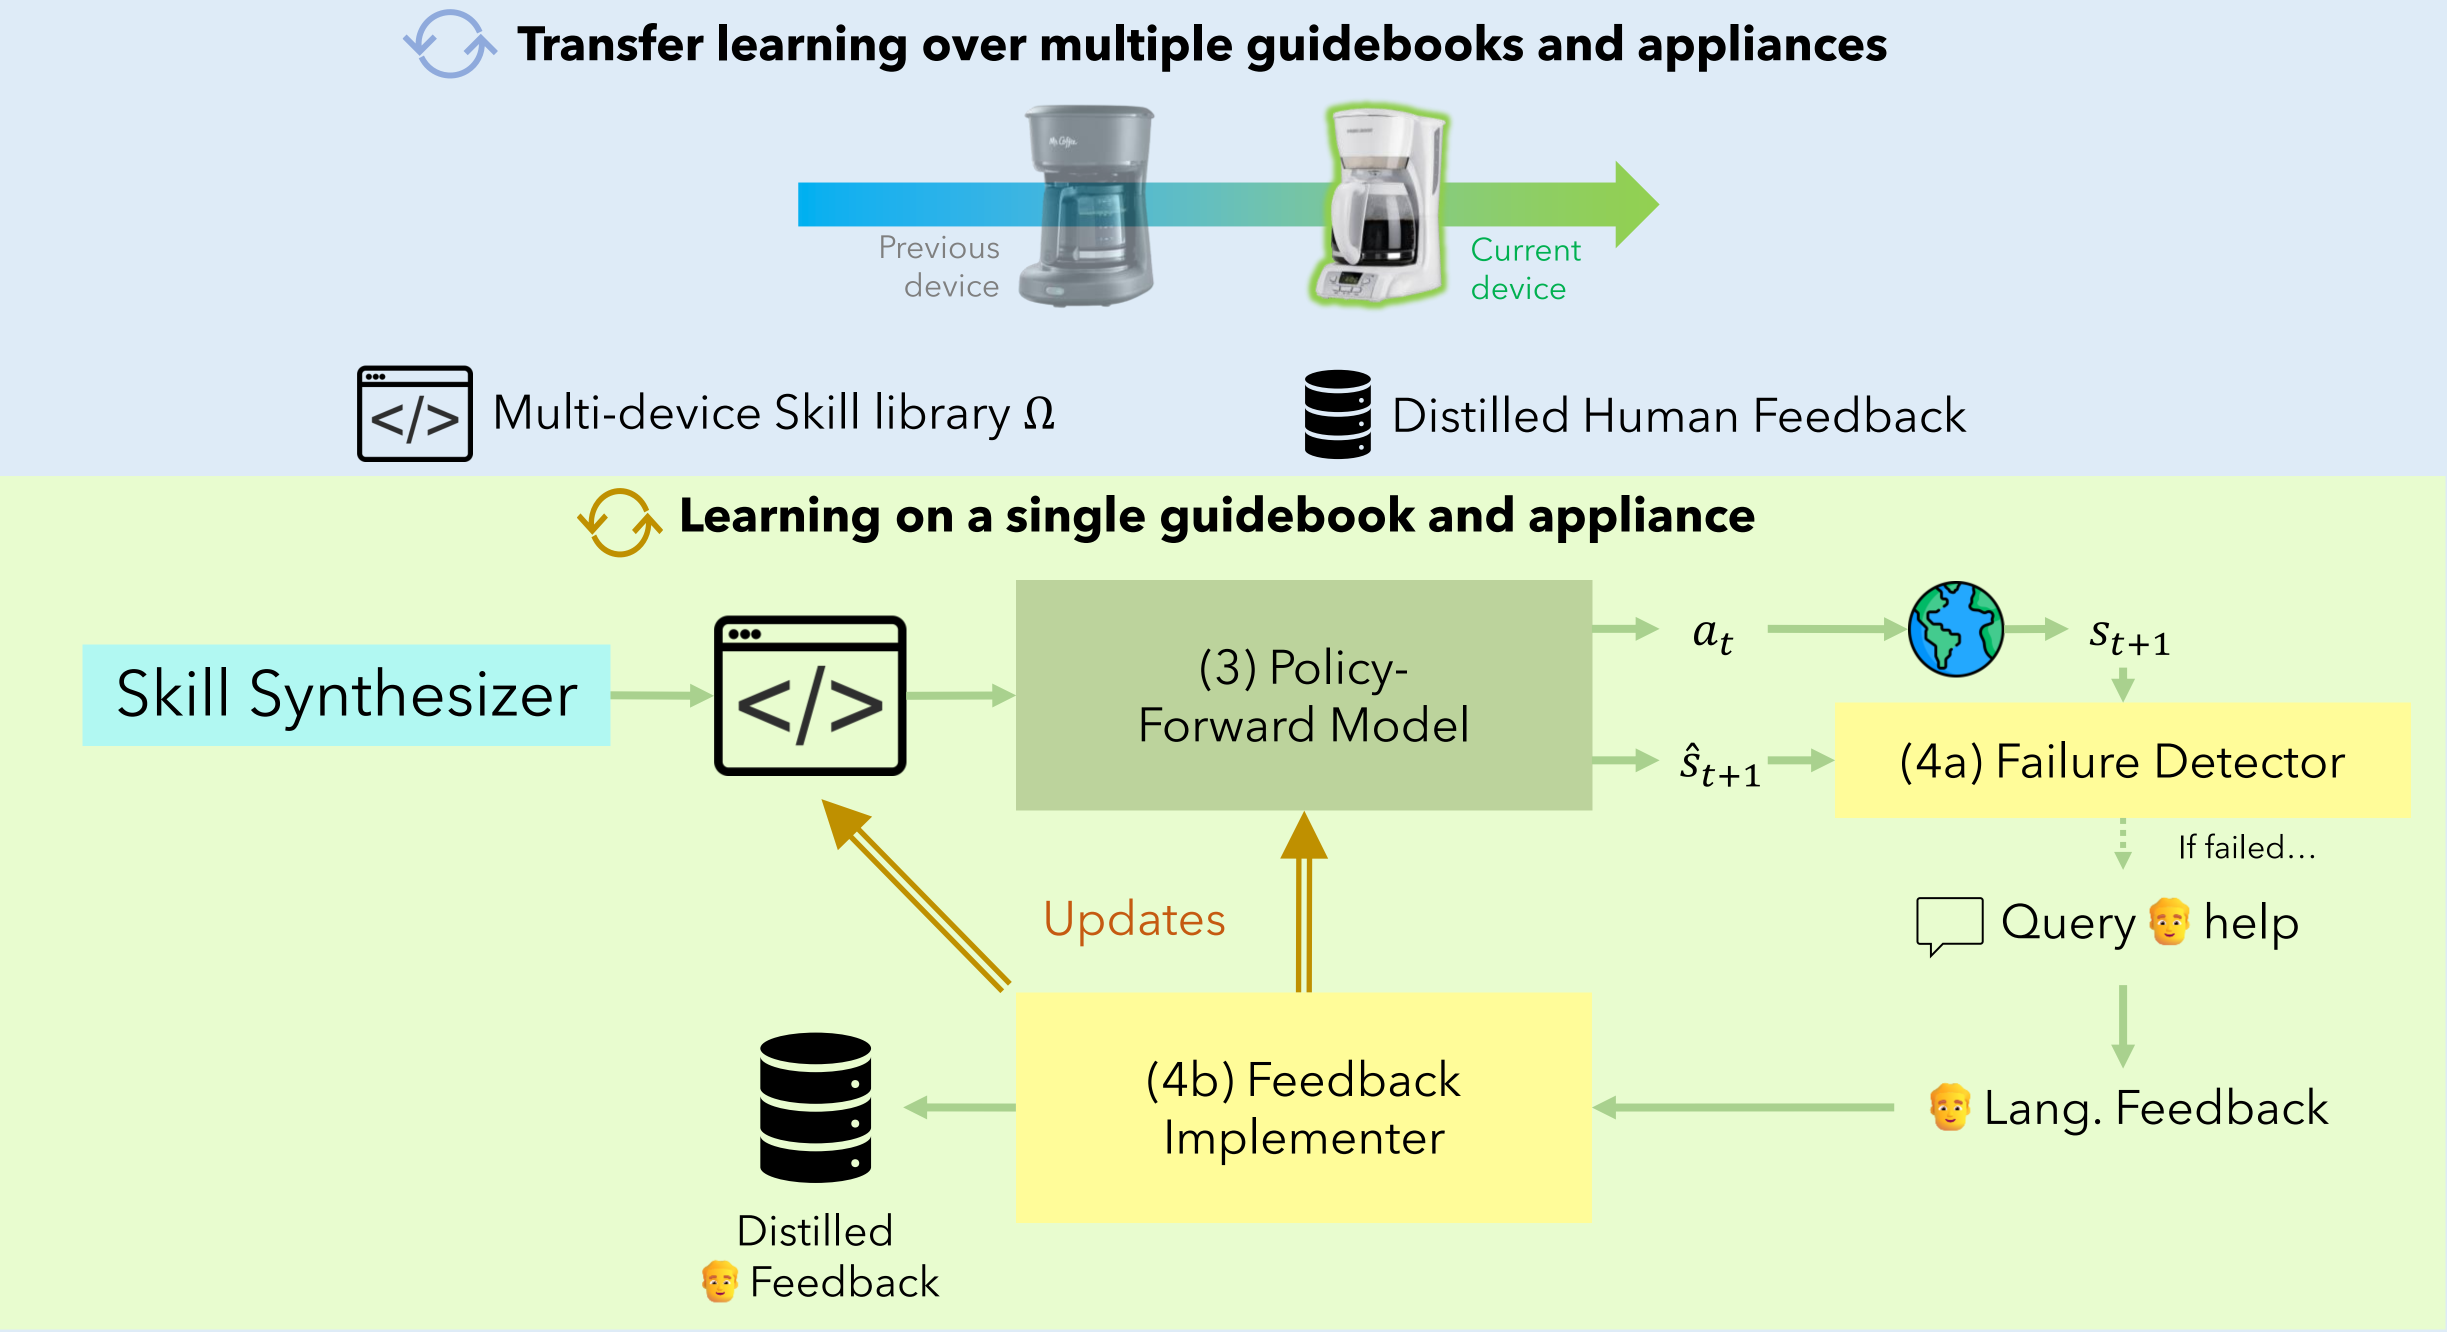
Transfer learning over multiple guidebooks and appliances
Previous device
Current device
Multi-device Skill library Ω
Distilled Human Feedback
Learning on a single guidebook and appliance
Skill Synthesizer
(3) Policy-Forward Model
a_t
s_{t+1}
$\hat{s}_{t+1}$
(4a) Failure Detector
If failed…
Query help
Updates
Lang. Feedback
(4b) Feedback Implementer
Distilled Feedback

Can BookWorm learn new devices faster?

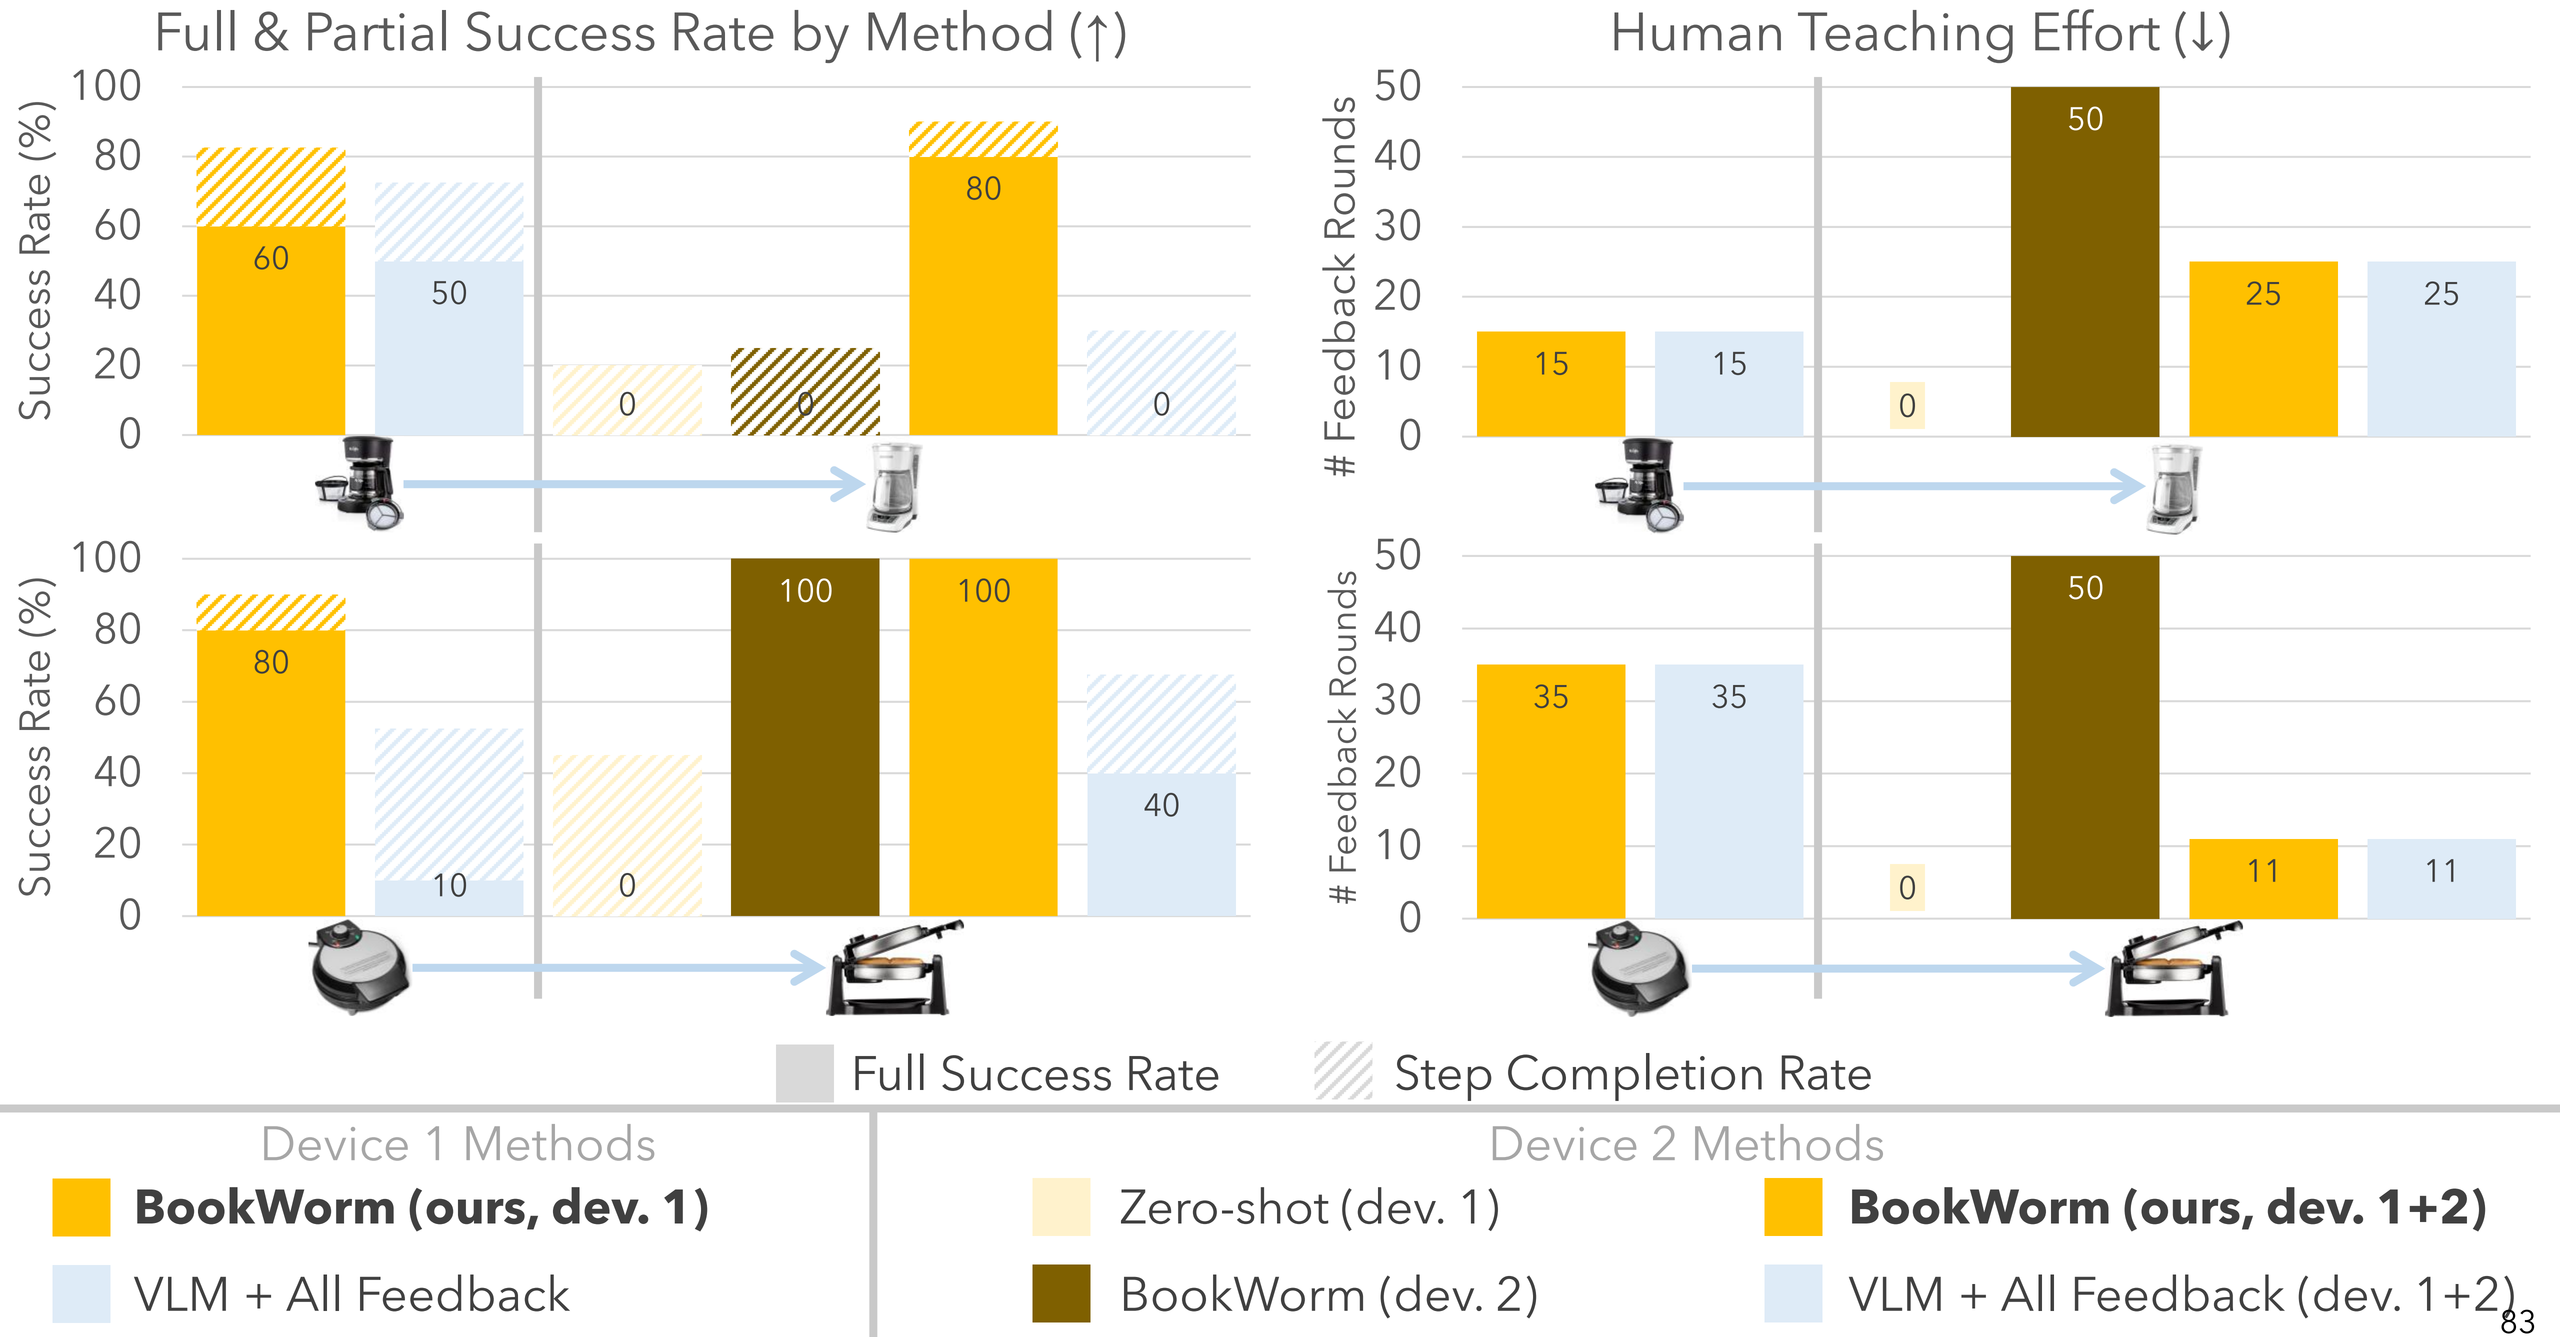

Can BookWorm learn new devices faster?

Full & Partial Success Rate by Method (↑)

Human Teaching Effort (↓)

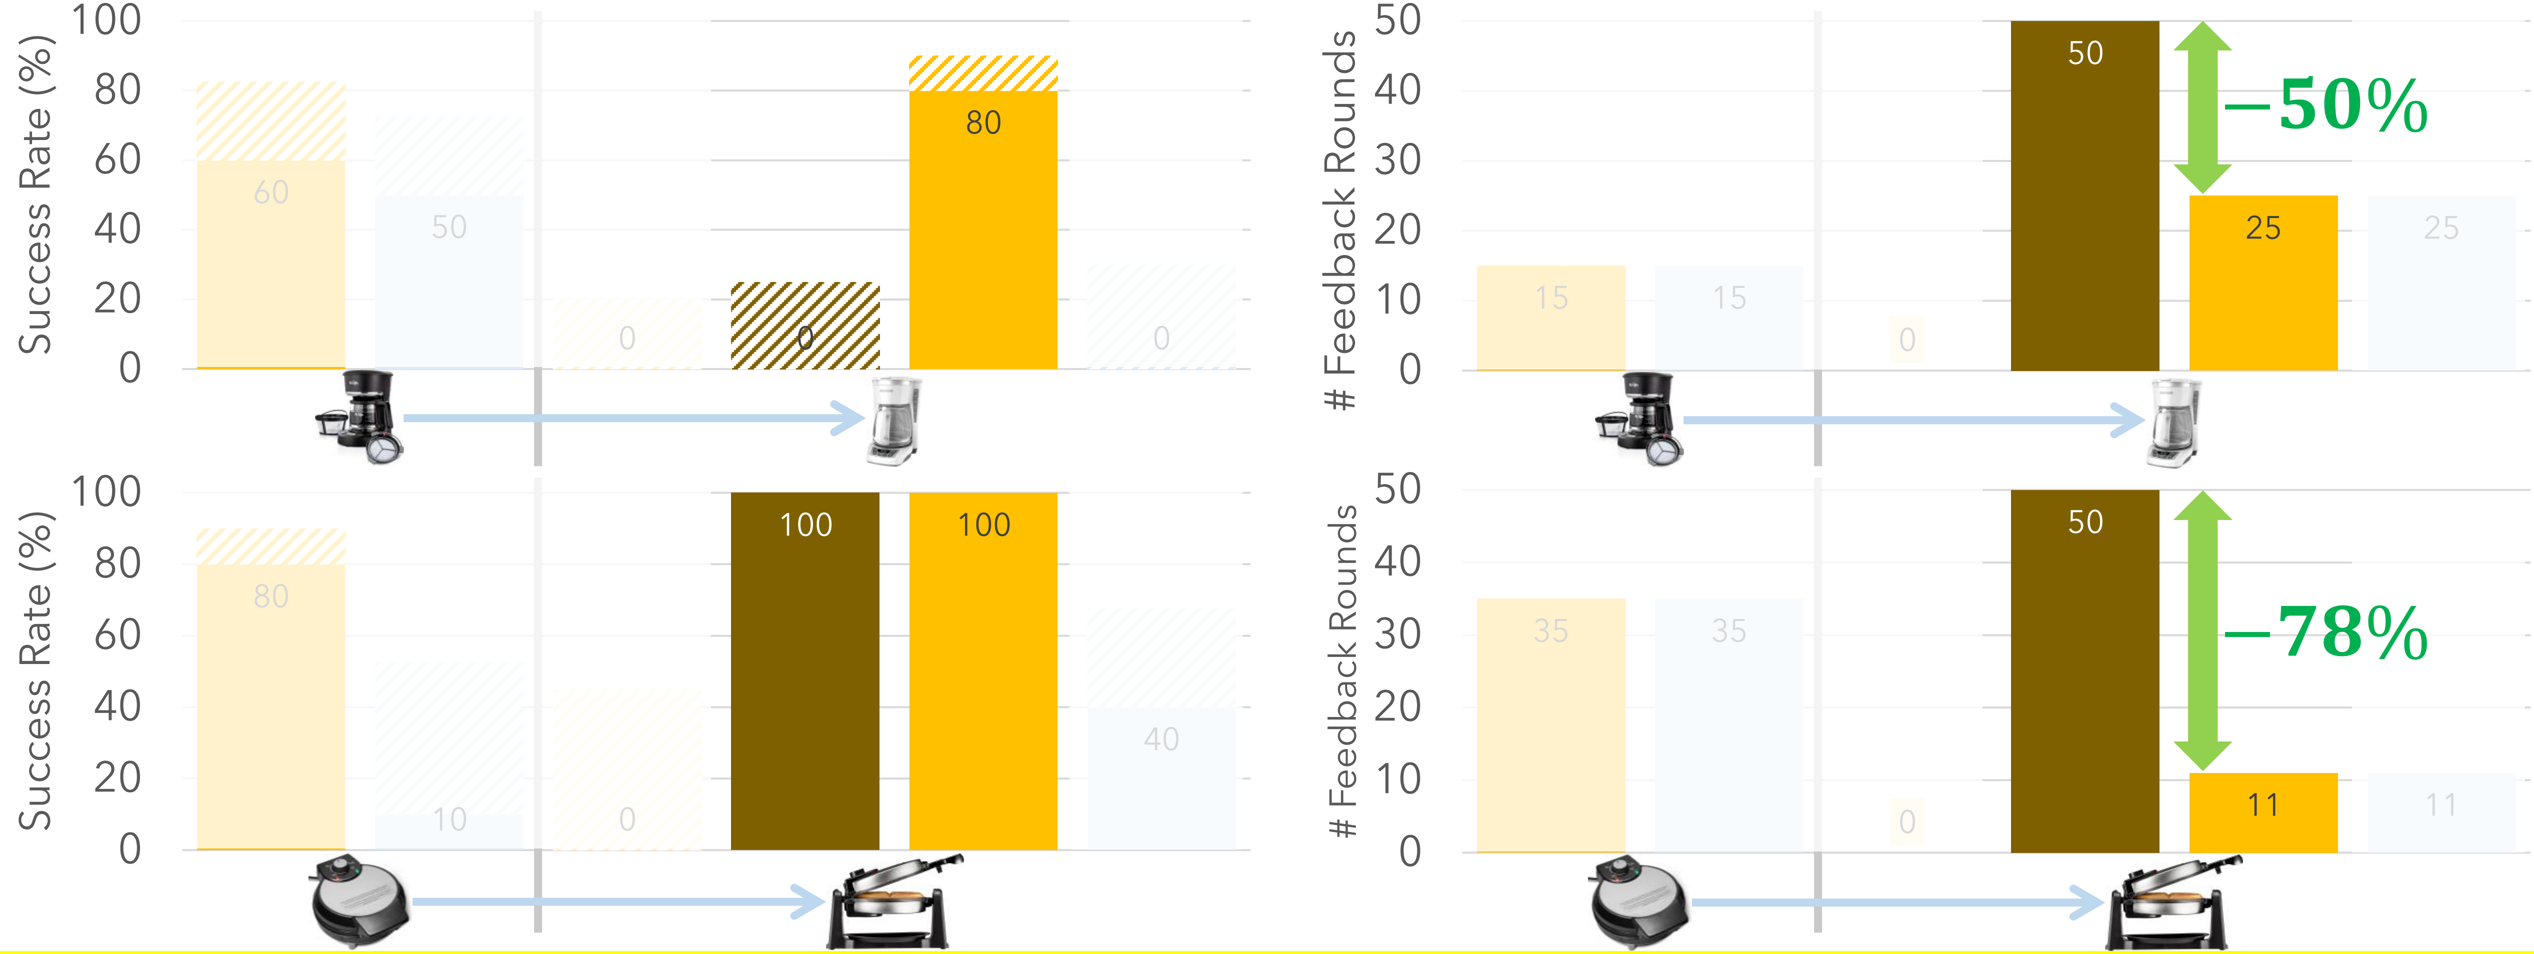

By leveraging prior device knowledge, BookWorm transfers efficiently to new (within-class) devices with ~60% less human feedback while matching/exceeding performance.

BookWorm (ours, dev. 1)

VLM + All Feedback

Zero-shot (dev. 1)

BookWorm (dev. 2)

BookWorm (ours, dev. 1+2)

VLM + All Feedback (dev. 1+2)


10x

Key Takeaways

Guidebooks provide action, dynamics, and entity grounding information, each crucial to acquiring a good manipulation prior.

BookWorm enables learning new within-class devices with ~60% less human effort.

BookWorm

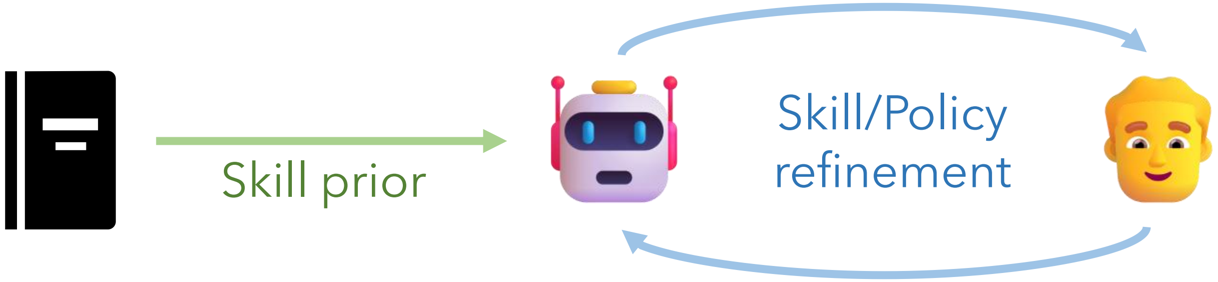

Outline

Language Improves Robotic Capabilities through…

Semantics for Generalization

1. DeL-TaCo: Simultaneous Learning from Demo + Language
2. Lang4Sim2Real: Language to Bridge the Sim2Real Gap

Communicative Medium

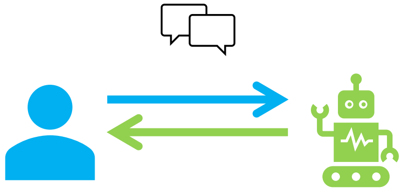

3. MicoBot: Bidirectional Dialog for Human-Robot Collaboration
4. BookWorm: Learning to operate new devices with guidebooks and human feedback

Future Work and Conclusion

Future Work: Semantics for Generalization

Building bridges with language to other modalities for better sample efficiency

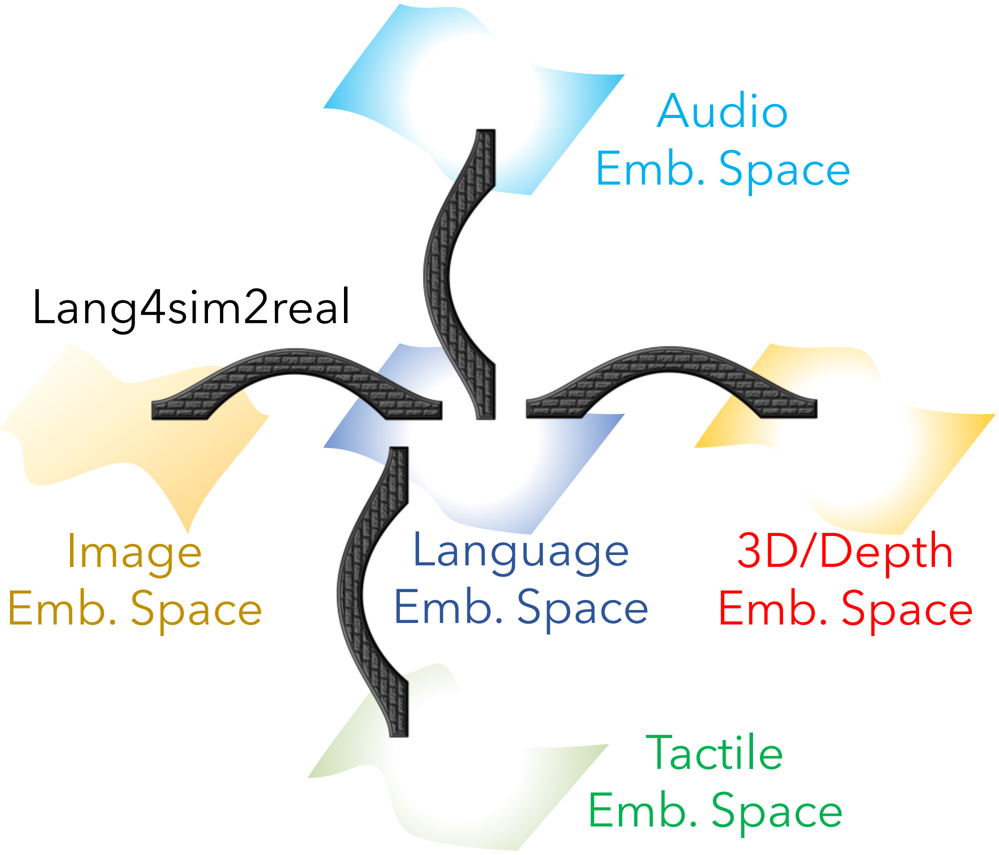

Future Work: – Communication

Going beyond task allocation dialog…

MicoBot: Mixed-initiative task allocation
+ DeL-TaCo: Mixed-initiative teaching

Mixed-initiative information sharing

- Keeping both agents safe

Future Work: 👨 – 🤖 Communication

Understanding **analogies** and **connections** to learn new tasks

👨 Add pepper to my steak.

I don't know how. 🤖

👨 The motion is like pouring and shaking the pepper simultaneously.

Conclusion

Language Augments Robotic Manipulation.

Semantics for Generalization

Enables generalization to new tasks (**DeL-TaCo**), domains (**Lang4sim2real**), and devices (**BookWorm**) with tens of demos and/or language instructions.

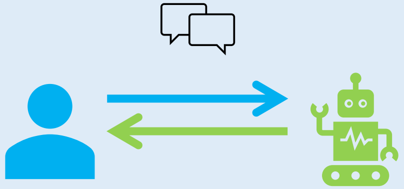

- Communication

Enables better collaboration with (**MiCoBot**) and faster learning from (**BookWorm**) diverse humans and human-written guidebooks.

What can we do next?

Build bridges with language to new modalities for sample-efficient robot learning

Richer communication for teaching and learning from robots

Thank you!

Advisors

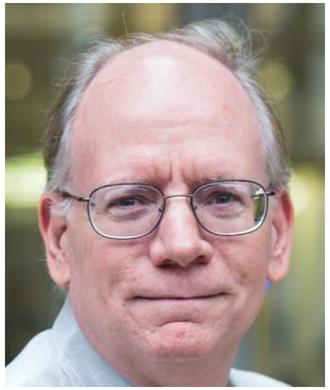

Prof. Ray Mooney

Prof. Roberto Martín-Martín

Committee Members

Prof. Peter Stone

Prof. Yonatan Bisk

Thank you!

Thesis Research Collaborators




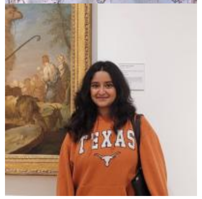

Chengshu Li
Adeline Foote
Will Reger
Cosmo Wu
Luca Macesanu
Arnav Balaji
Ruchira Ray
Da Liu




More Research Collaborators during PhD

Rutav Shah
Yifeng Zhu
Mina Huh
Huihan Liu
Jonathan Yang
Homer Walke
Avi Singh
Jedrzej Orbik
Zhanpeng He
Sergey Levine
Jiajun Wu
Yuke Zhu
Amy Pavel
Maya Cakmak
Jane Shi
Fan Wang
Michael Schultz

Brad Knox
Shasank Govindarajan
Natalia Cislo
Mary Thompson
Michael Baldea
Bruce Eldridge
Ayub Lakhani
Mark Nixon
Claire Wright
Beth Hood
Erik Frey
Baruch Tabanpour
Shuozhe Li
Steven Shi
Soroush Nasiriany

Thank you!

LUNAR and RobIn Labmates,
For a very supportive, open-minded, thought-provoking atmosphere for research.

Thank you!

Finally…

Teachers and Mentors
Friends
My Family

Back to Conclusion