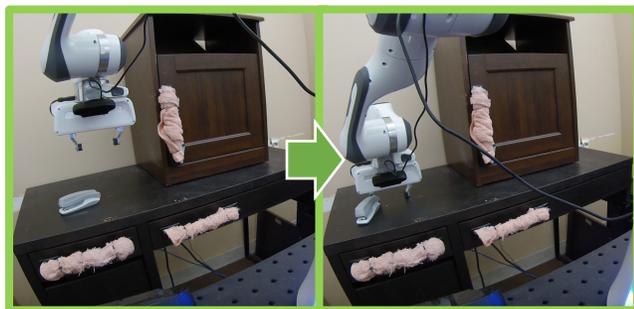


Learning Language-Conditioned Robot Behavior from Offline Data and Crowd-Sourced Annotation

Suraj Nair¹, Eric Mitchell¹, Kevin Chen¹, Brian Ichter², Silvio Savarese¹, Chelsea Finn^{1,2}

¹Stanford University, ²Robotics at Google

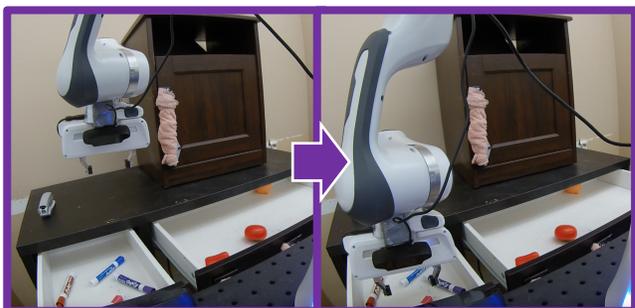
“Move the stapler”



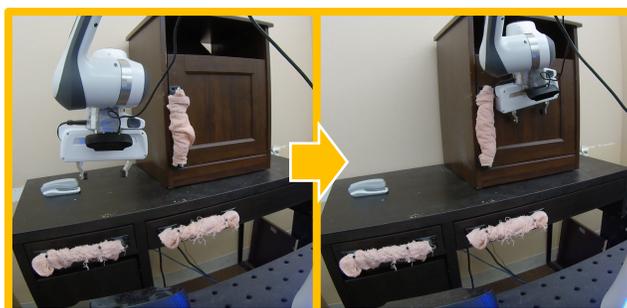
“Open the left drawer”



“Reach the marker”



“Reach the cabinet”

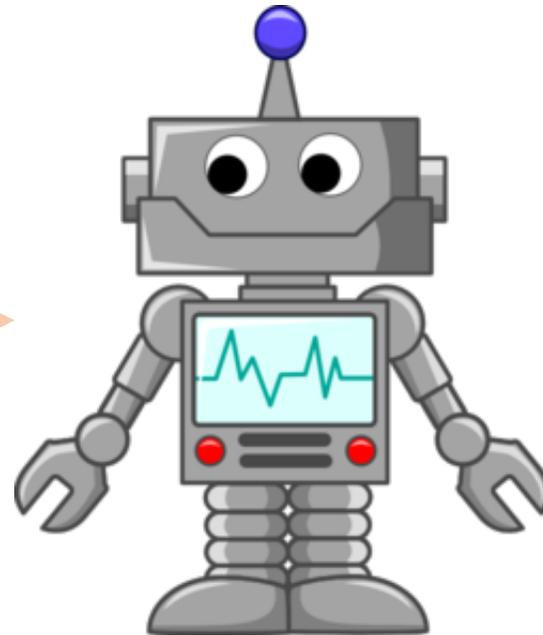
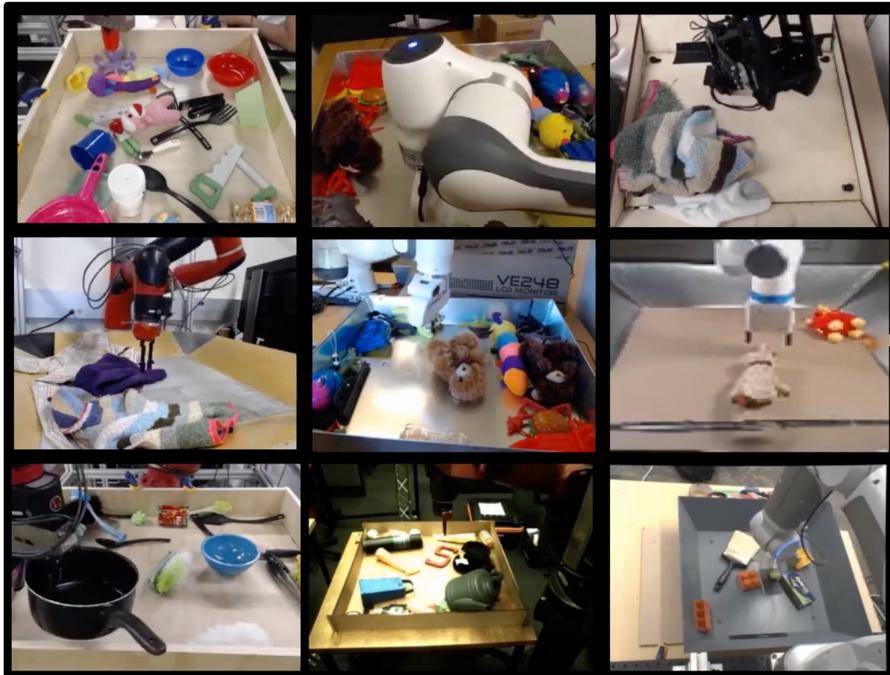


“Open the right drawer”



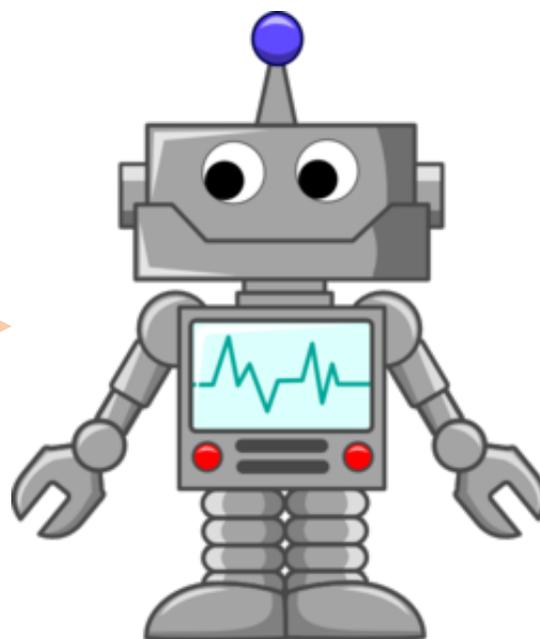
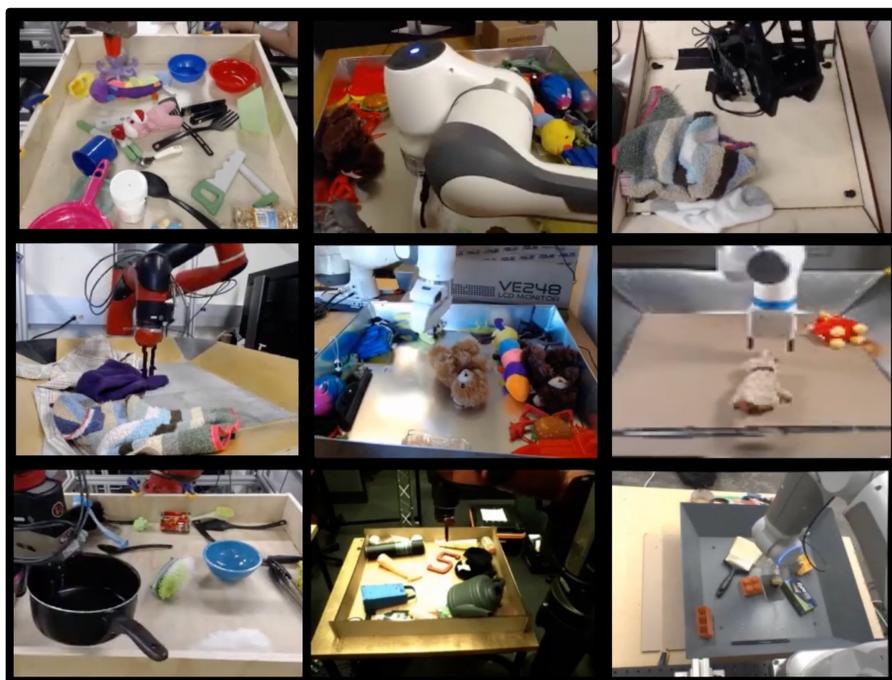
Multi-Task Learning from Offline Data

Offline Data



Multi-Task Learning from Offline Data

Offline Data



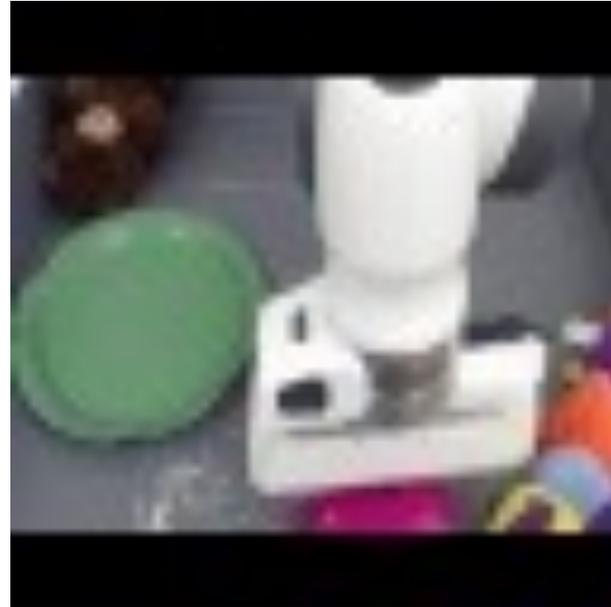
Humans need easy and effective ways of specifying tasks

Goal Image Task Specification

Initial Image



Goal Image

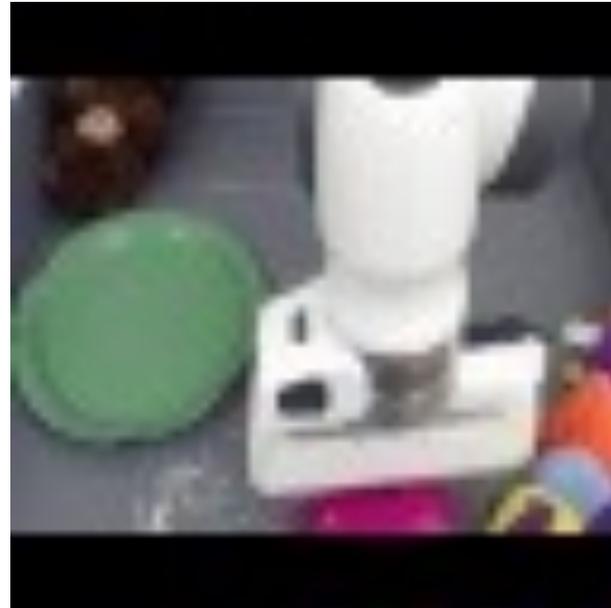


Goal Image Task Specification

Initial Image



Goal Image



Requires human effort

Goal Image Task Specification

Initial Image



Goal Image



Requires human effort
Task over-specification

Goal Image Task Specification

Initial Image



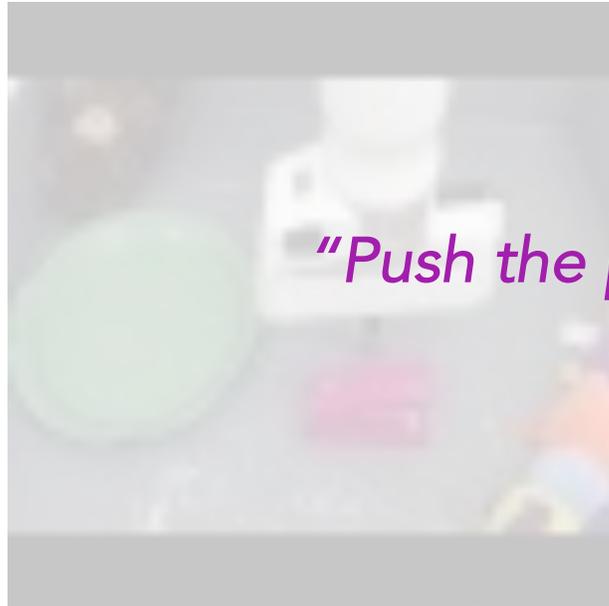
Goal Image



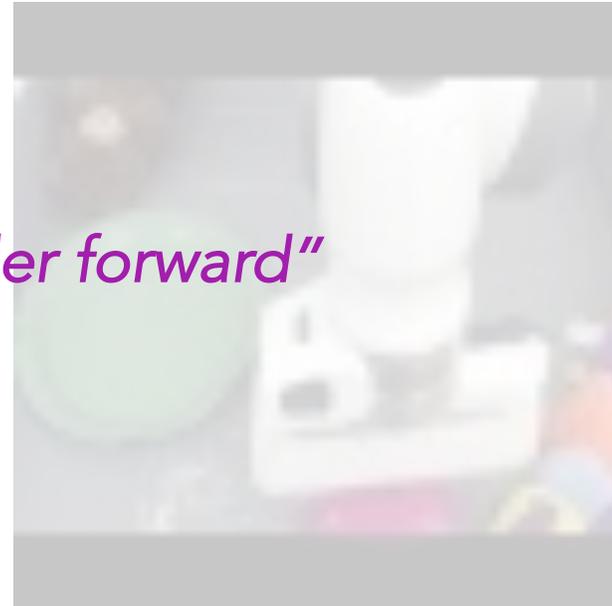
Requires human effort
Task over-specification
Cannot specify persistent behavior

Natural Language Task Specification

Initial Image



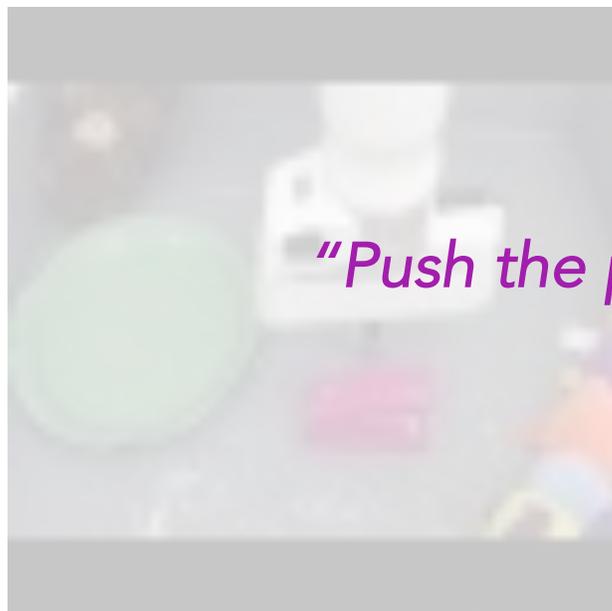
Goal Image



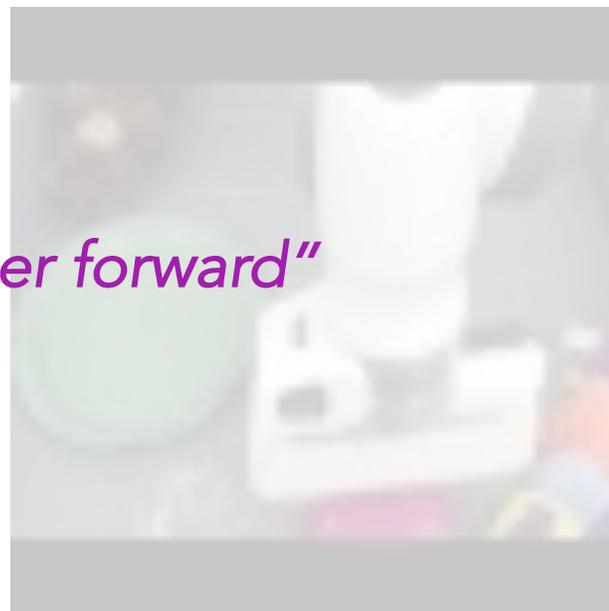
"Push the pink stapler forward"

Natural Language Task Specification

Initial Image



Goal Image



"Push the pink stapler forward"

Easy for humans to provide
Can flexibly represent tasks

Our Objective

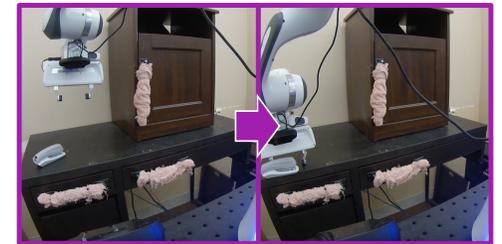
Learn *language-conditioned visuomotor manipulation skills* from:
offline datasets +
crowd-sourced annotation

Language-Conditioned Task Completion

"Move the stapler"



"Push the small gray stapler around on top of the black desk"



"Open the left drawer"



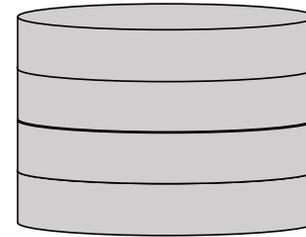
"Open the right drawer"



Our Objective

Learn *language-conditioned visuomotor manipulation skills* from:
offline datasets +
crowd-sourced annotation

Offline Robot
Data

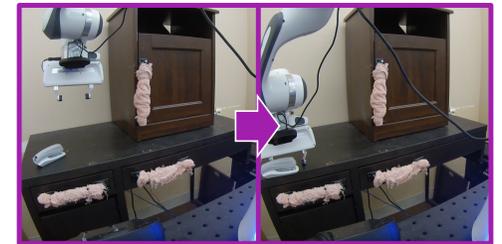


Language-Conditioned Task Completion

“Move the stapler”



“Push the small gray stapler
around on top of the black desk”



“Open the left drawer”



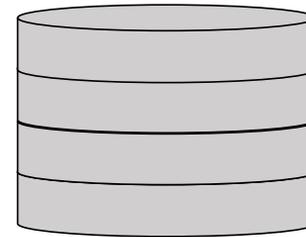
“Open the right drawer”



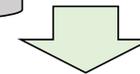
Our Objective

Learn *language-conditioned visuomotor manipulation skills* from:
offline datasets +
crowd-sourced annotation

Offline Robot
Data



Crowd-Sourced
Annotation

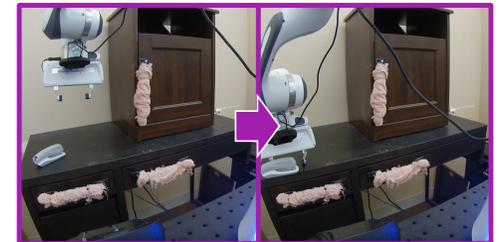


Language-Conditioned Task Completion

“Move the stapler”



“Push the small gray stapler
around on top of the black desk”



“Open the left drawer”



“Open the right drawer”



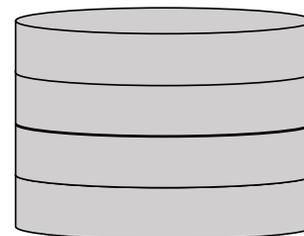
Our Objective

Learn *language-conditioned visuomotor manipulation skills* from:
offline datasets +
crowd-sourced annotation

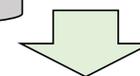
Offline data:

- Highly suboptimal
- From autonomous exploration /existing replay buffers

Offline Robot Data



Crowd-Sourced Annotation



Language-Conditioned Task Completion

“Move the stapler”



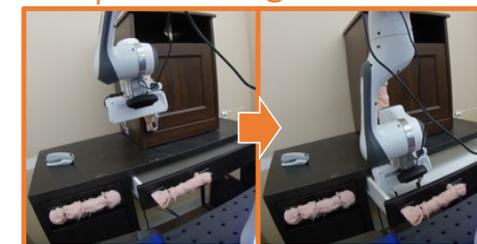
“Push the small gray stapler around on top of the black desk”



“Open the left drawer”



“Open the right drawer”



Our Objective

Learn *language-conditioned visuomotor manipulation skills* from:
offline datasets +
crowd-sourced annotation

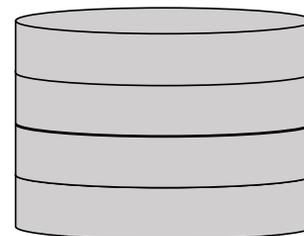
Offline data:

- Highly suboptimal
- From autonomous exploration /existing replay buffers

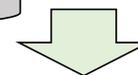
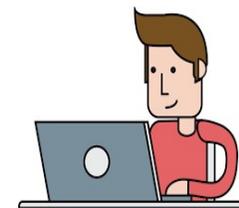
Crowd-sourced annotation:

- Describes what task (if any) is completed

Offline Robot Data



Crowd-Sourced Annotation

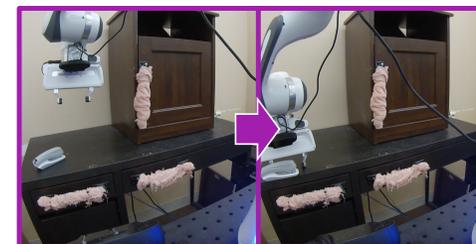


Language-Conditioned Task Completion

"Move the stapler"



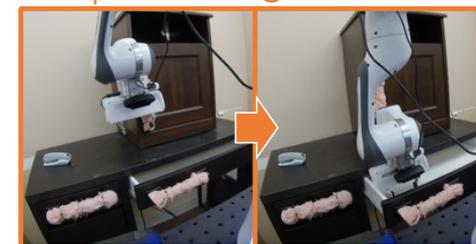
"Push the small gray stapler around on top of the black desk"



"Open the left drawer"



"Open the right drawer"



Key Idea

Actions in data cannot be treated as optimal

Key Idea

Actions in data cannot be treated as optimal

start → *finish* of episode
completes the annotated instruction

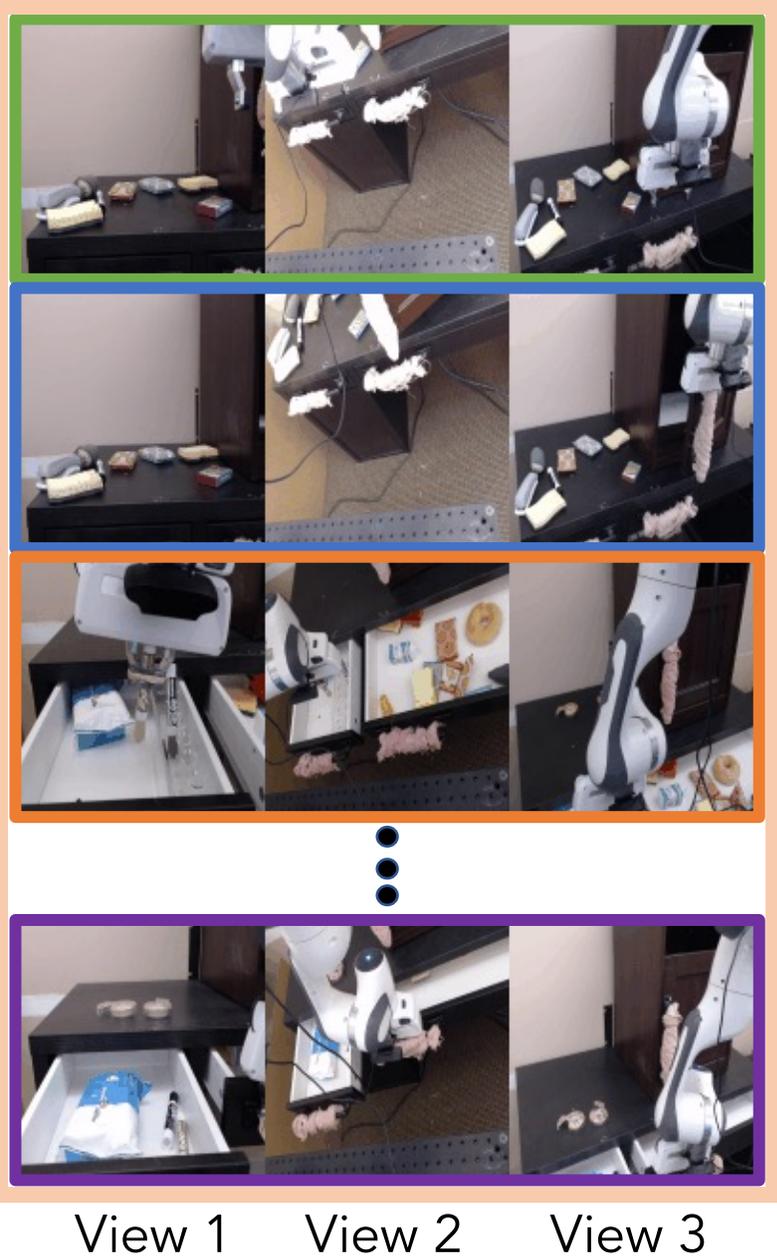
Key Idea

Actions in data cannot be treated as optimal

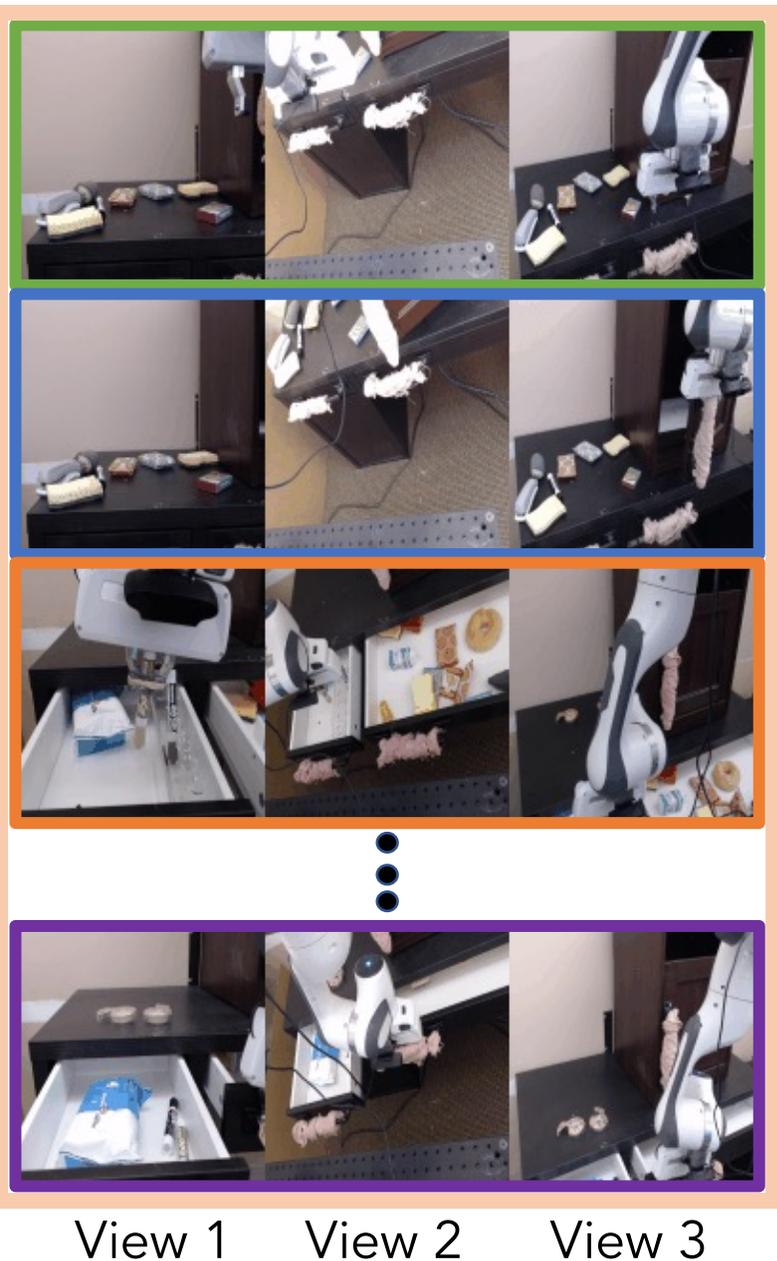
start → *finish* of episode
completes the annotated instruction

Learn *language-conditioned reward function* for offline RL

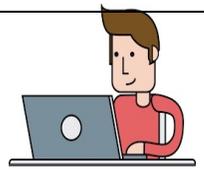
Offline Data



Offline Data

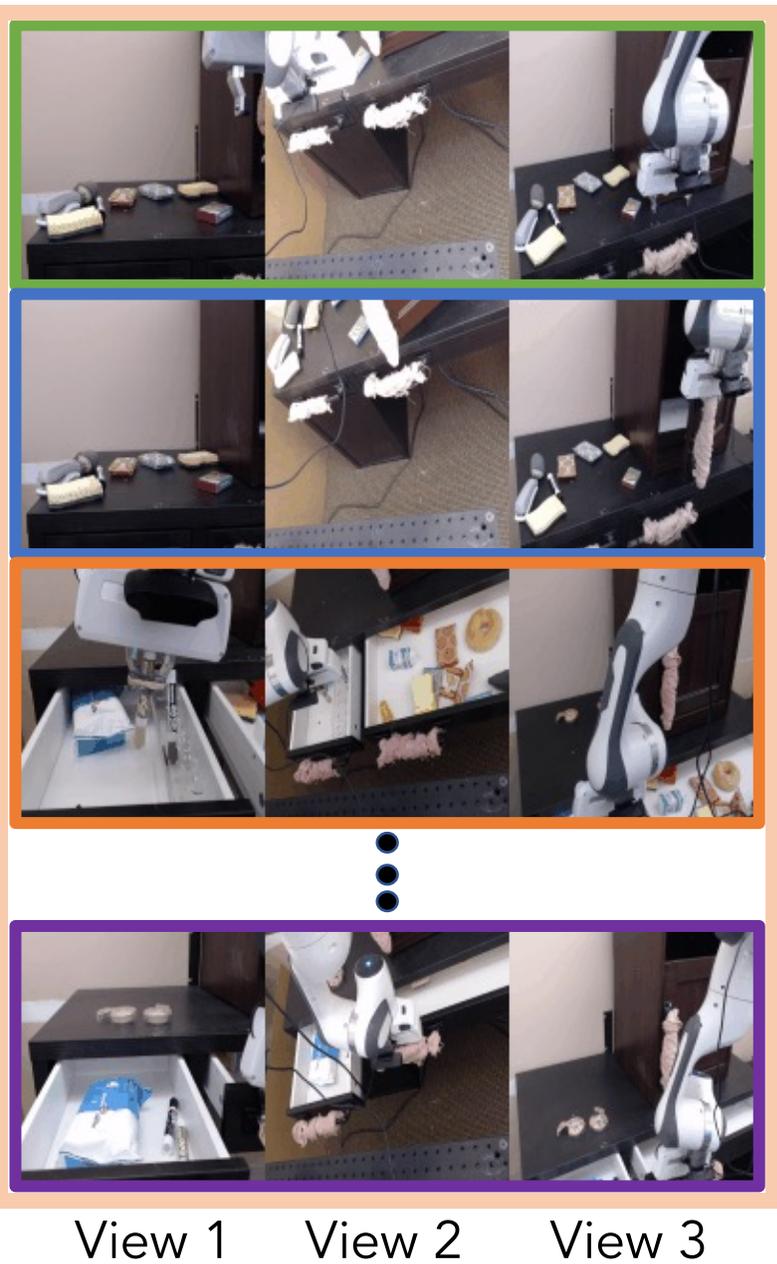


“Do nothing”
“Open the cabinet on the desk”
“Try to insert the marker into the hole”
...
“Open the right drawer halfway”



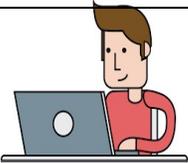
Crowdsourced Annotation

Offline Data



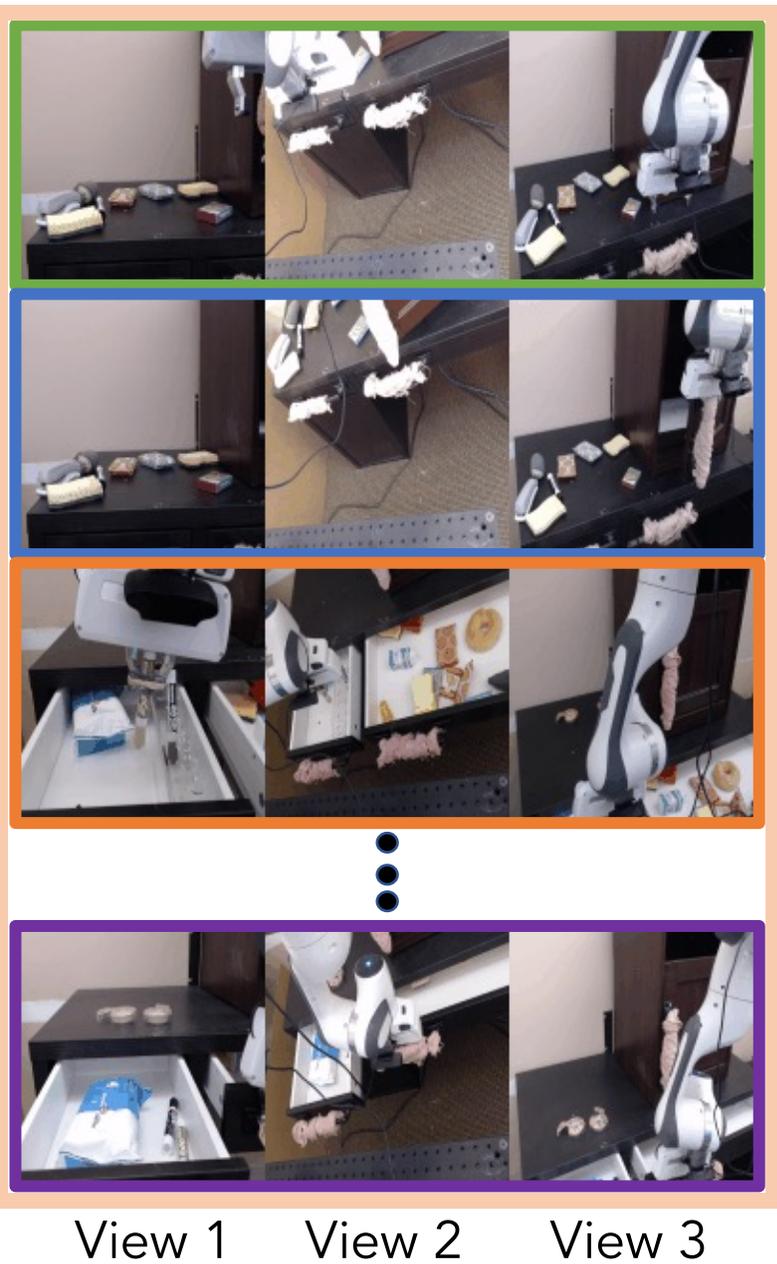
Language-Conditioned
Reward Learning

“Do nothing”
“Open the cabinet on the desk”
“Try to insert the marker into the hole”
...
“Open the right drawer halfway”



Crowdsourced Annotation

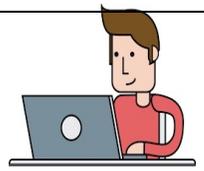
Offline Data



Visual Dynamics Learning

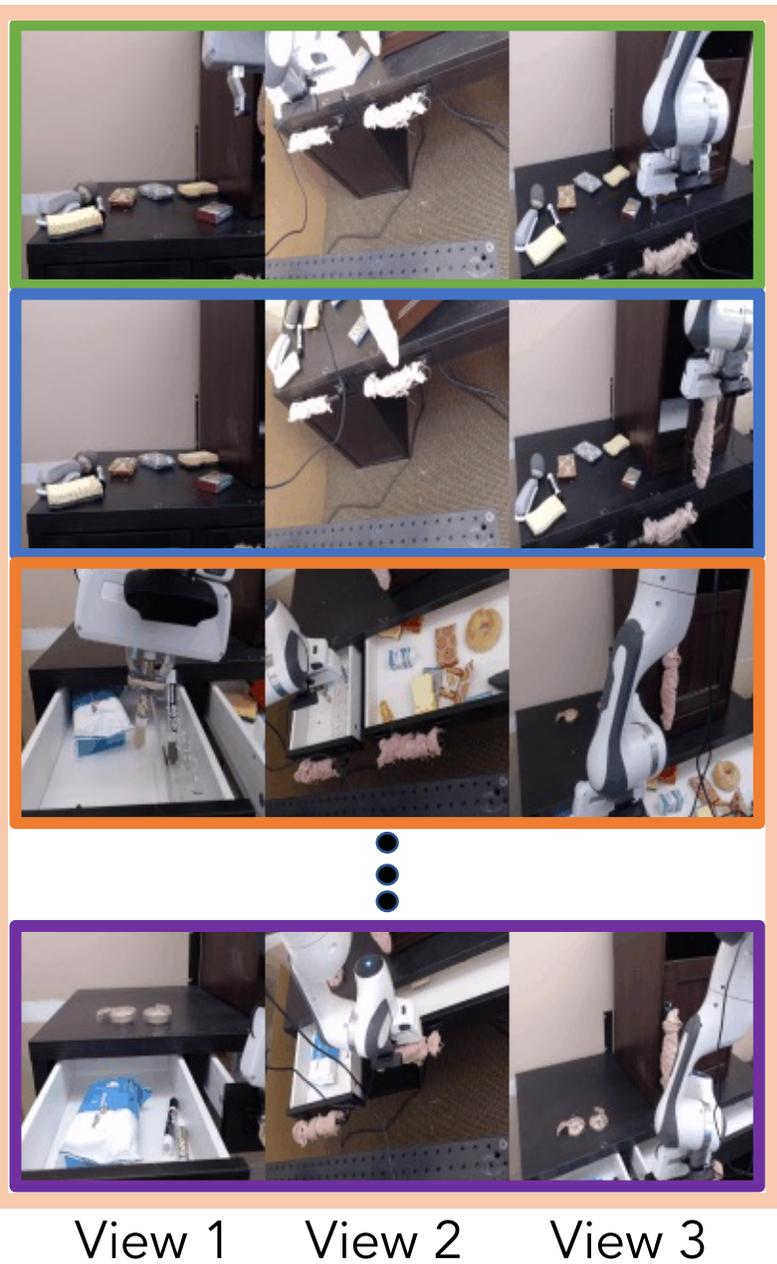
Language-Conditioned
Reward Learning

“Do nothing”
“Open the cabinet on the desk”
“Try to insert the marker into the hole”
...
“Open the right drawer halfway”



Crowdsourced Annotation

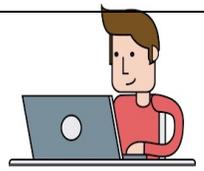
Offline Data



Visual Dynamics Learning

Language-Conditioned Reward Learning

"Do nothing"
"Open the cabinet on the desk"
"Try to insert the marker into the hole"
...
"Open the right drawer halfway"

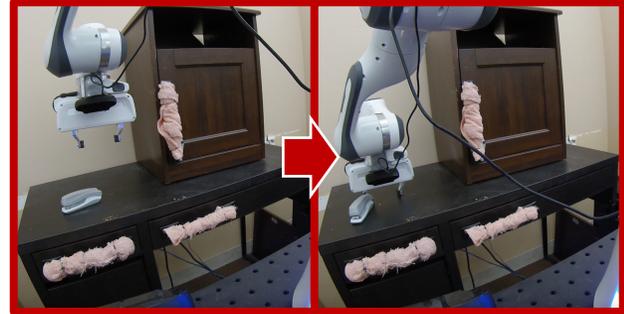


Crowdsourced Annotation

Evaluation:

Model Predictive Control

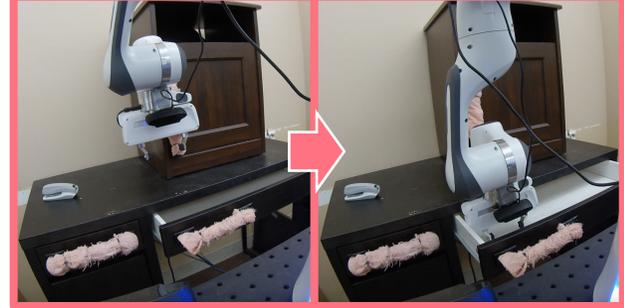
"Move the stapler"



"Open the left drawer"



"Open the right drawer"

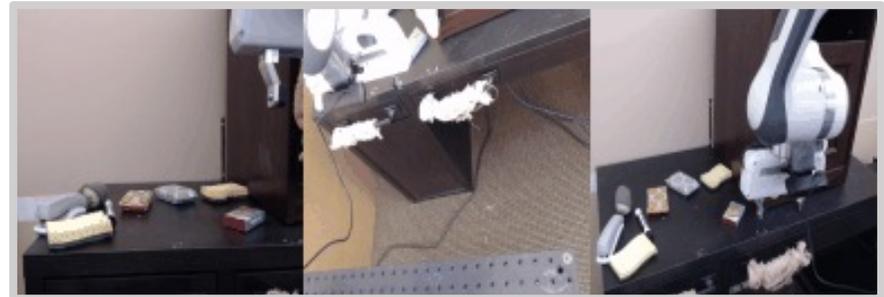


Language-Conditioned Offline Reward Learning (LOReL)

Learn binary classifier of instruction completion

Instruction: "Open the cabinet on the desk"

S_0



S_T



LOReL

Label = 1

Language-Conditioned Offline Reward Learning (LOReL)

Learn binary classifier of instruction completion

Positives:

- Episode and corresponding annotations

Instruction: "Open the cabinet on the desk"

S_0



S_T



LOReL

Label = 1

Language-Conditioned Offline Reward Learning (LOReL)

Learn binary classifier of instruction completion

Negatives:

- Different episodes

Instruction: "Open the cabinet on the desk"



LOReL

Label = 0

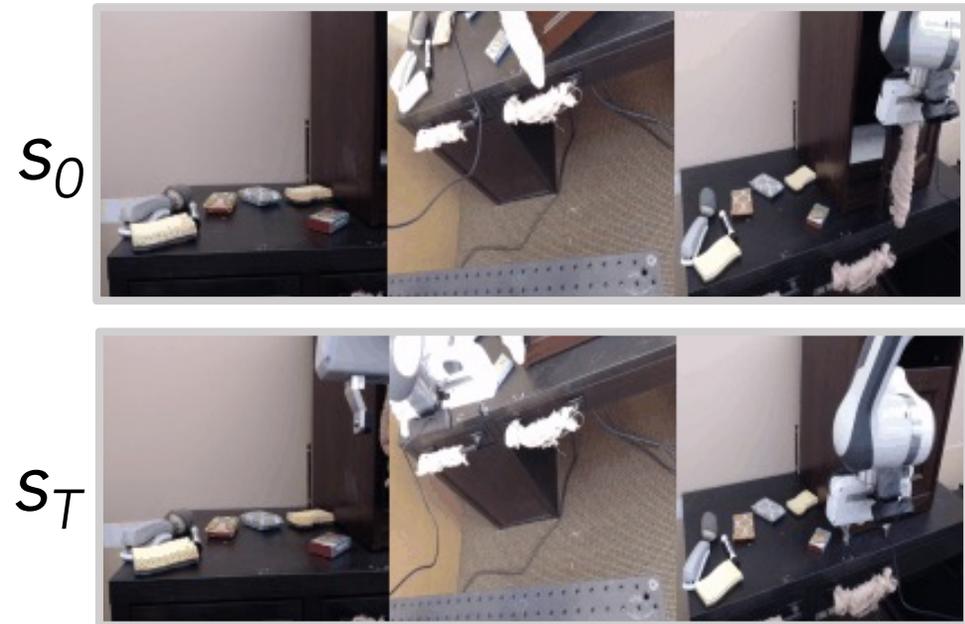
Language-Conditioned Offline Reward Learning (LOReL)

Learn binary classifier of instruction completion

Negatives:

- Different episodes
- Episode reversed

Instruction: "Open the cabinet on the desk"



LOReL

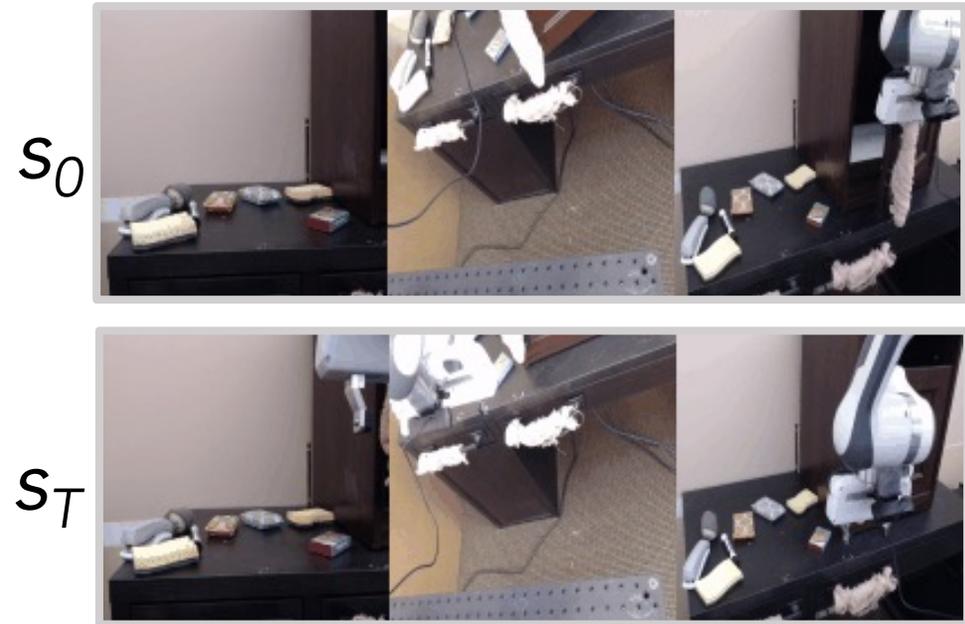
Label = 0

Language-Conditioned Offline Reward Learning (LOReL)

Learn binary classifier of instruction completion

Pretrained language models to enable generalization from limited data

Instruction: "Open the cabinet on the desk"



LOReL

Label = 0

Language-Conditioned Offline Reward Learning (LOReL)

Initial State

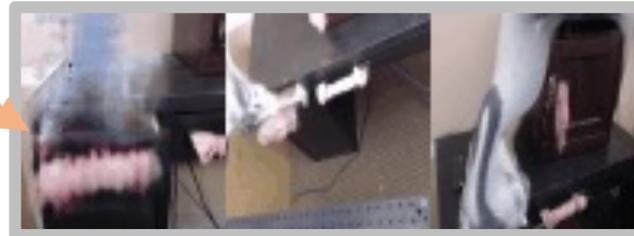


View 1 View 2 View 3

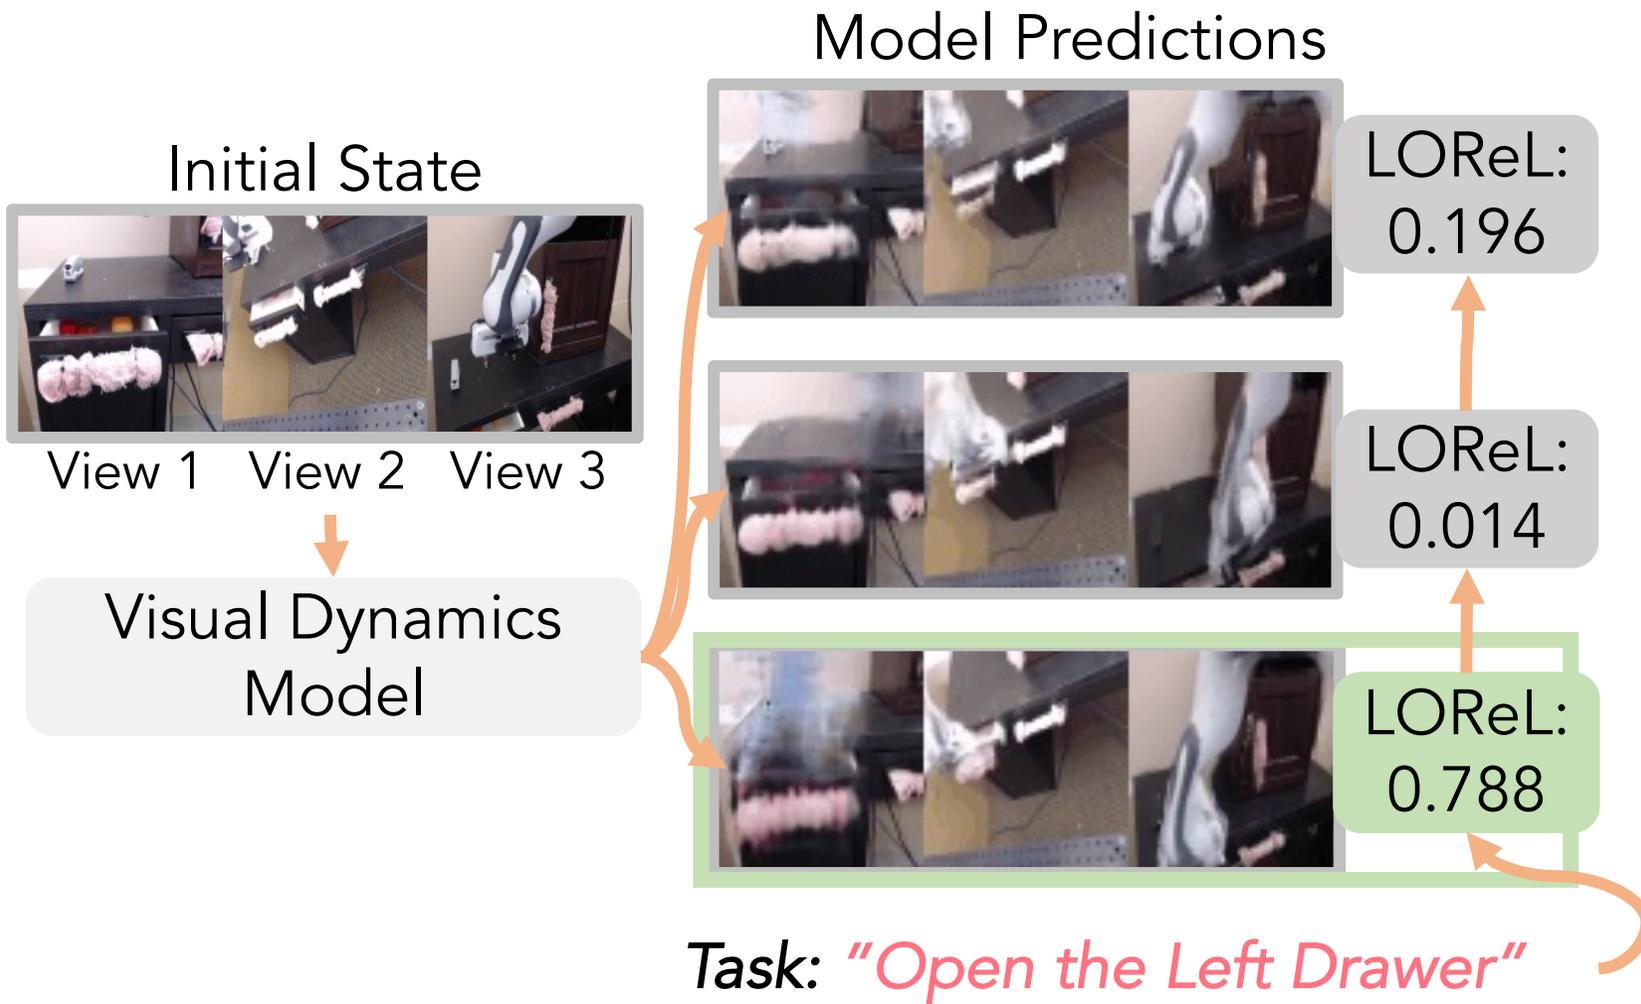


Visual Dynamics Model

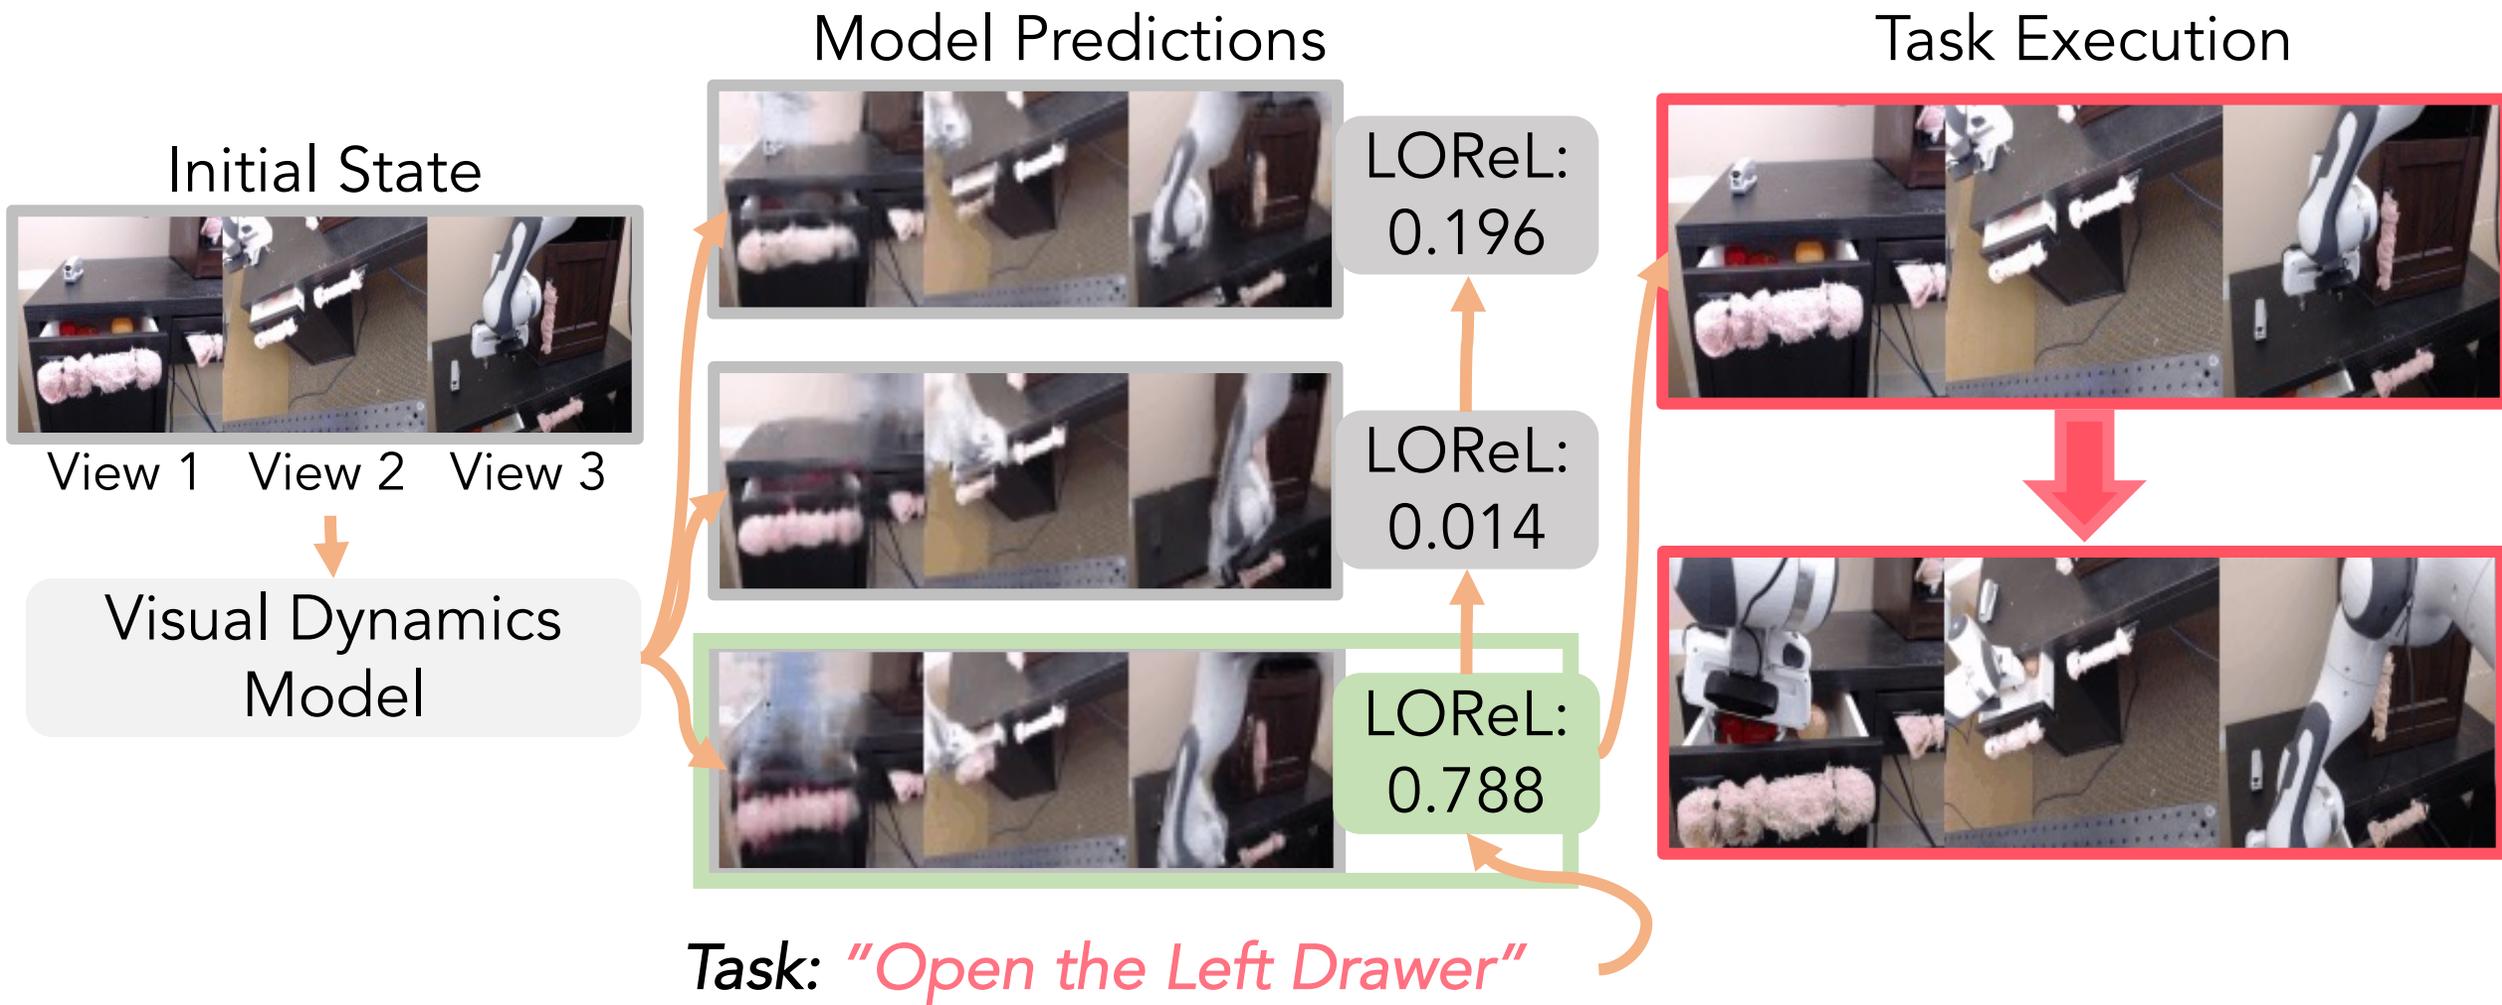
Model Predictions



Language-Conditioned Offline Reward Learning (LOReL)

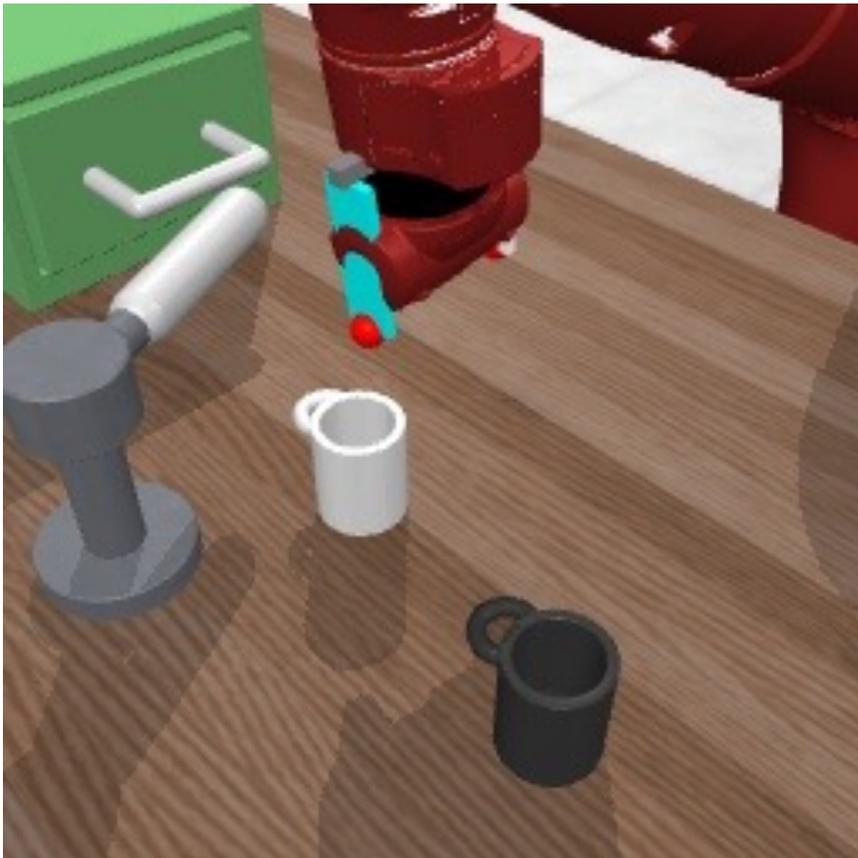


Language-Conditioned Offline Reward Learning (LOReL)



Experiment 1: How does LOReL enable effective language-conditioned behavior compared to prior methods?

Simulated Environment

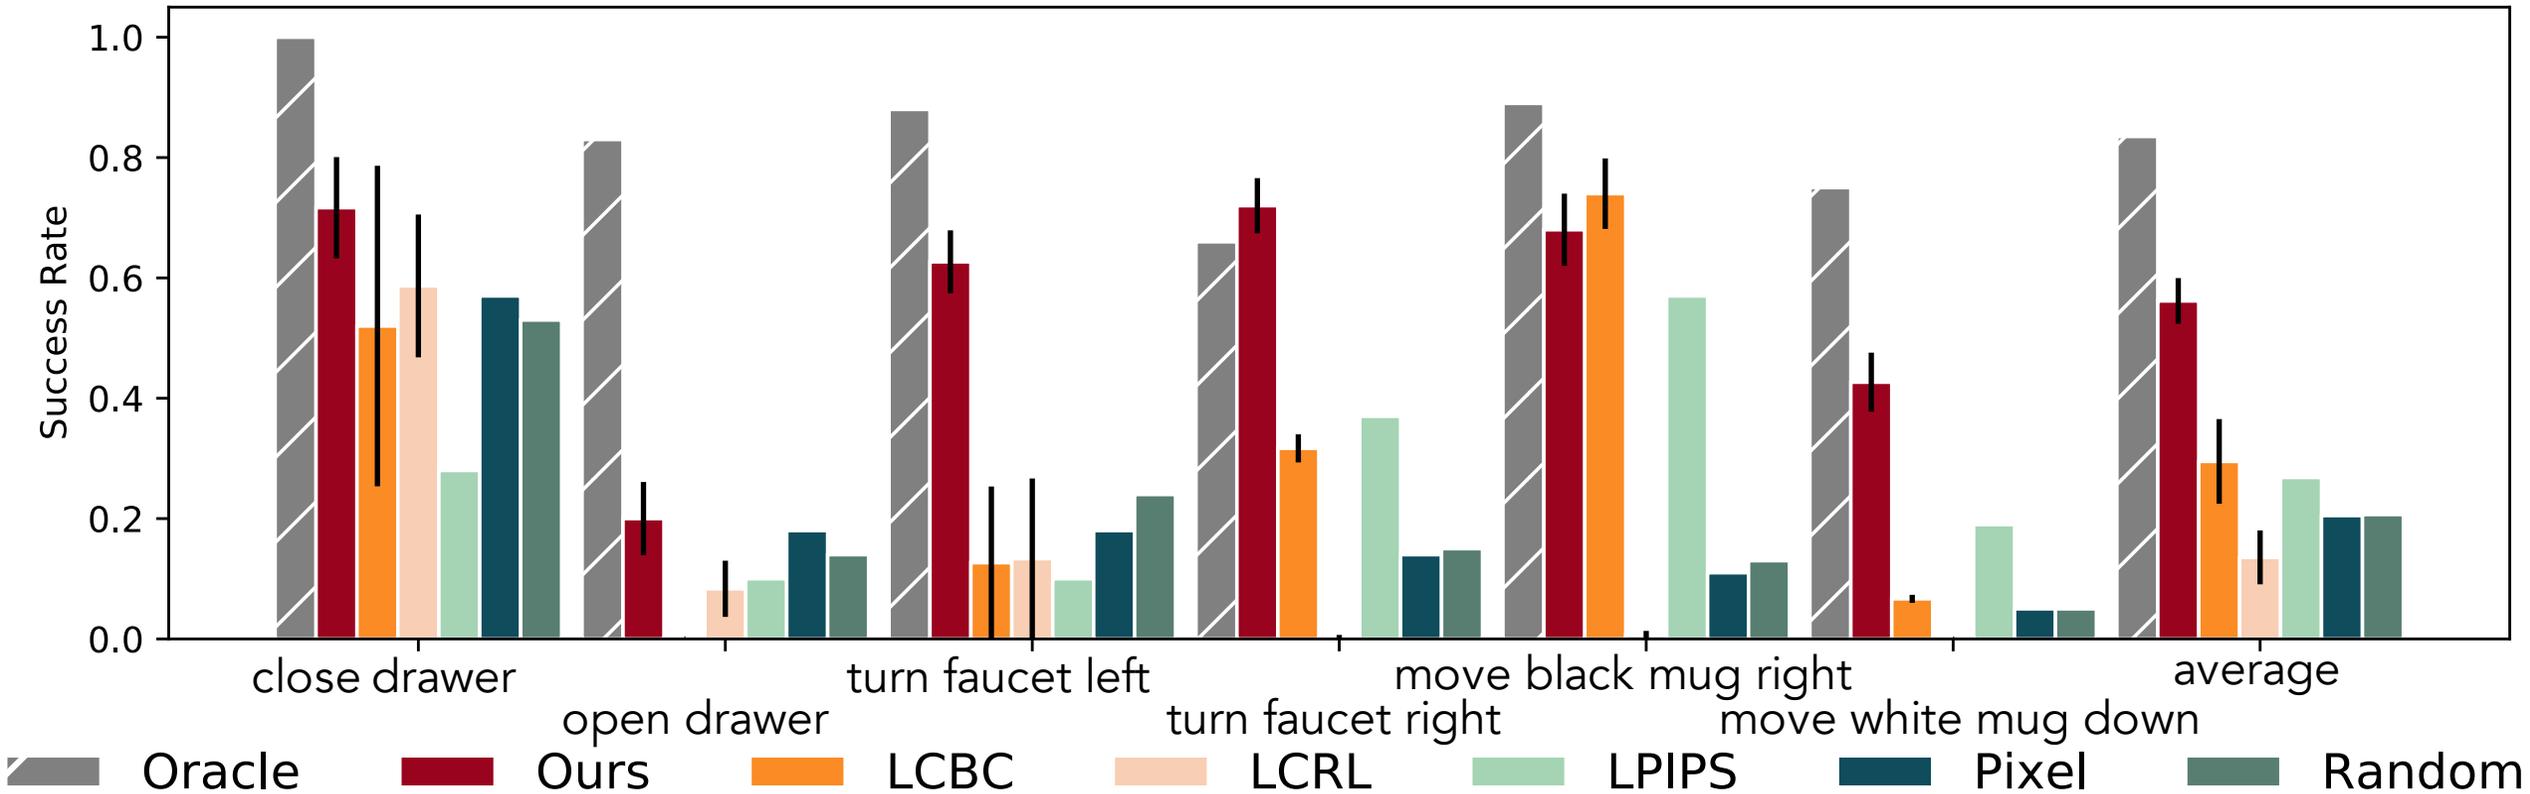


Procedurally Generated Annotations

- Close the drawer
- Move the white mug up and turn the faucet left
- Open the drawer and turn the faucet right
- Do nothing
- Move the white mug right
- Move the black mug down and left

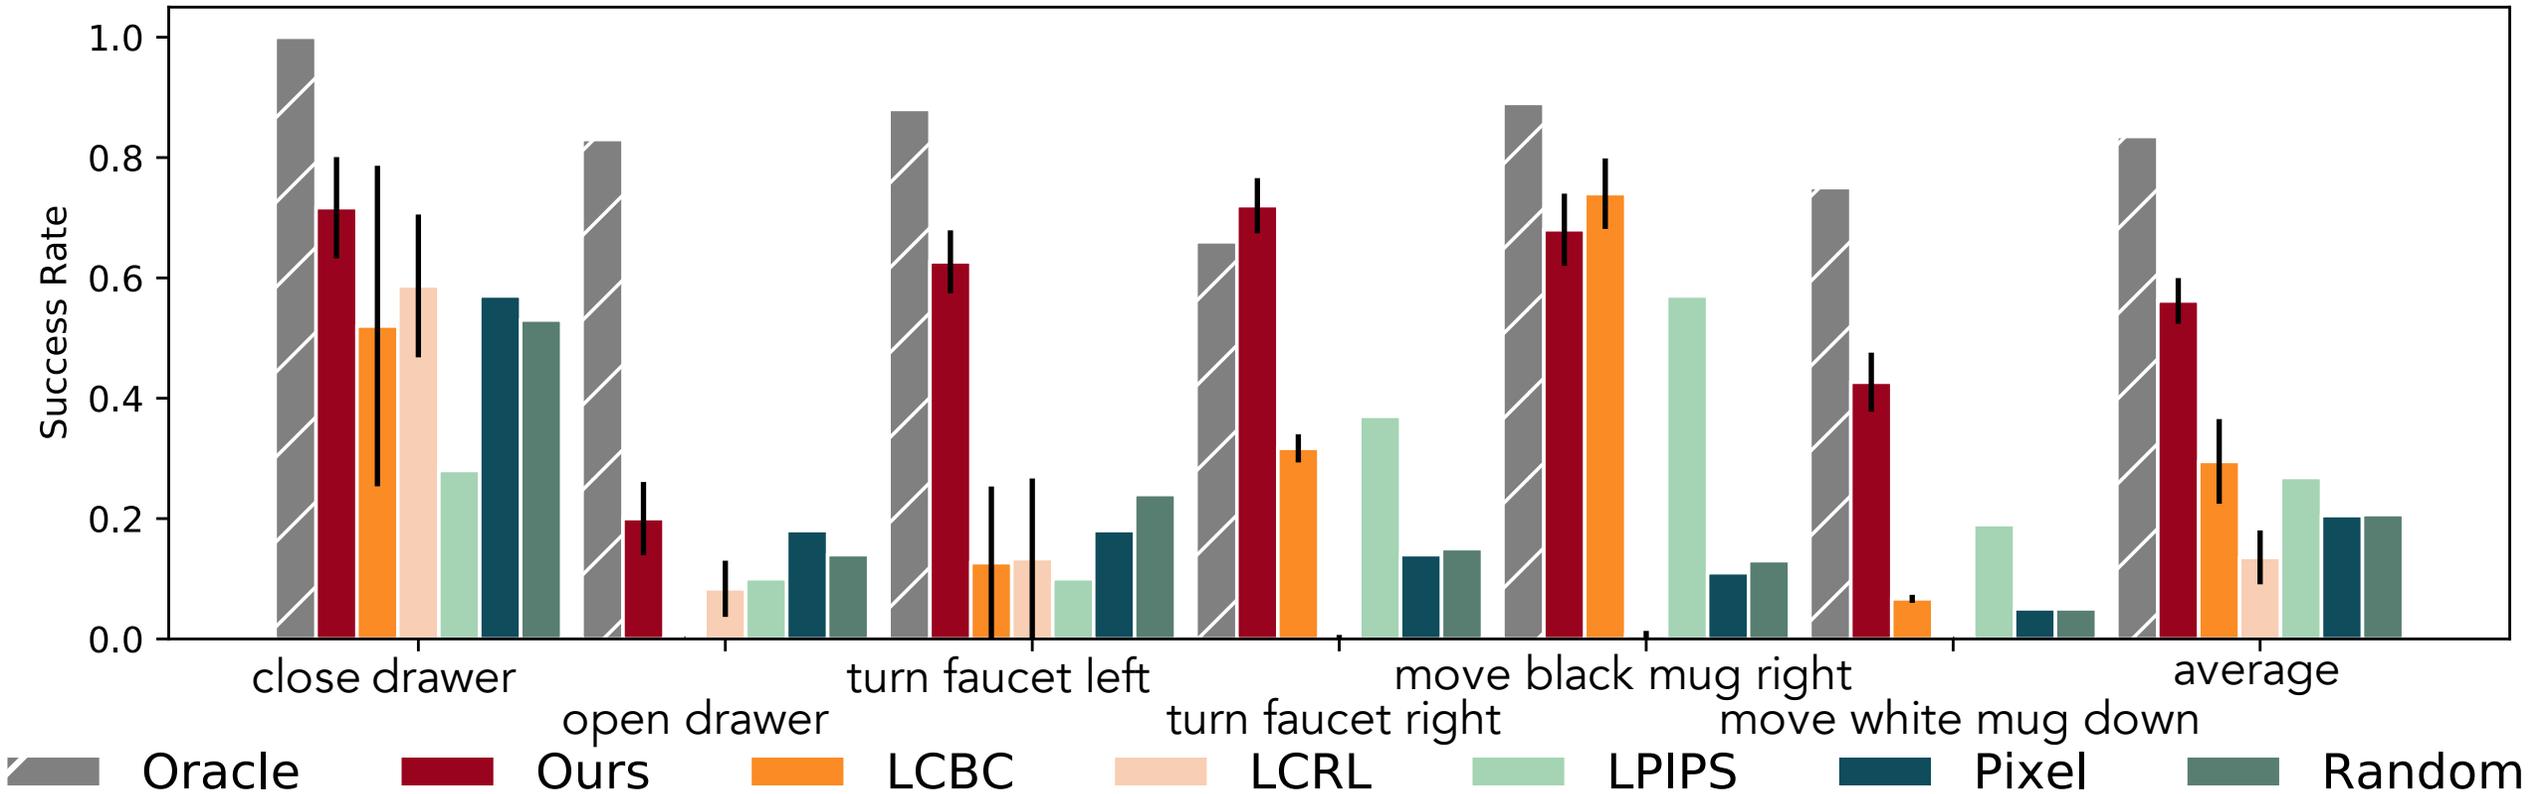


Experiment 1: How does LOReL enable effective language-conditioned behavior compared to prior methods?



- **LOReL** outperforms *language conditioned imitation learning/RL* by 25+%

Experiment 1: How does LOReL enable effective language-conditioned behavior compared to prior methods?



- **LOReL** outperforms *language conditioned imitation learning/RL* by 25+%
- **LOReL** outperforms goal image *Pixel/LPIPS* specification by 29+%

Experiment 2:

Can LOReL generalize to unseen natural language commands?

Success Rate	Ours	Ours (-PM)
Original	56 ± 1%	40 ± 1%
Unseen Verb	51 ± 3%	33 ± 2%
Unseen Noun	51 ± 1%	39 ± 4%
Unseen Verb + Noun	47 ± 2%	17 ± 3%
Unseen Natural Language	46 ± 2%	21 ± 1%

- LOReL can generalize to unseen instructions effectively (<10% drop in performance)
- Leveraging a pretrained model is essential for this generalization

Instruction	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6
Seen	Close drawer	Open drawer	Turn faucet left	Turn faucet right	Move black mug right	Move white mug down
Unseen Verb	shut drawer	pull drawer	rotate faucet left	rotate faucet right	push black mug right	push white mug down
Unseen Noun	close container	open container	turn tap left	turn tap right	move dark cup right	move light cup down
Unseen Verb+ Noun	shut container	pull container	rotate tap left	rotate tap right	push dark cup right	push light cup down
Human Provided	Push the drawer shut	Pull the drawer open	Rotate the tap counterclock wise	Rotate tap clockwise	Translate the black cup to the right	Translate the white cup down
	Push the drawer	Pull the handle	Turn faucet away from camera	Turn faucet towards camera	Move black mug away from drawer	Move with mug closer to the faucet
	Shut the drawer	Pull the drawer handle	Rotate nozzle left	Rotate nozzle right	Push black cup right	Bring white cup down
	shut drawer	Pull the drawer open	faucet counterclock wise	faucet clockwise	black mug right	white mug down
	Slide the drawer closed	Pull open the drawer	Rotate the faucet left	Rotate the faucet right	Slide the black mug right	Push the white mug down and left
	Shut the drawer	open the dresser	Turn the faucet to the left	Turn the faucet to the right	Move the dark mug to the right	Move the lighter mug down
	shut the dresser	Pull the drawer.	Turn the faucet to the left	rotate handle rightward	push black cup right	shift white mug down
	Shut the drawer.	Unclose the cabinet	rotate handle to the left	Turn faucet clockwise.	Move black mug right.	Pull white mug to the front.
	Shut the cupboard		Turn faucet counter clockwise.	Twirl valve right	Shift dark cup right	Reposition white glass down
			Spin nozzle left			

Experiment 3: Is LOReL effective on a real robot?

Robot Domain



3000 episodes taken from
concurrent work performing
online RL

Replay buffer consists of
successful and unsuccessful
attempts of many skills

2 annotations per episodes
collected using AMT

Experiment 3: Is LOReL effective on a real robot?

Robot Domain



Instructions: Given a video of a robot, write a sentence summarizing its behavior.

Write the sentence in the form of a command (i.e. "pick up the stapler" rather than "picking up the stapler").

Note: you can see the robot's behavior from 3 different camera viewpoints.

Some examples:

"Open the left drawer" **NOT** "Opening the left drawer",

"Reach the marker" **NOT** "Reaching the marker",

"Push the stapler up" **NOT** "Pushing the stapler up",

"Insert the plug into the socket"

"Open the cabinet"

Only write down an instruction that the robot successfully completed.

For example, if the robot tries to pick up the marker and gets close but fails, you should put "Reach to the marker" or "Grasp next to the marker" rather than "pick up the marker".

If the robot does not successfully complete any task, write something like:

"Do nothing"

"Wave the arm in the air"

If you can't see the robot or can't tell what the robot is doing, write:

"NA"



Write a command which describes the robot's behavior in the video...

Submit

Experiment 3: Is LOReL effective on a real robot?

Robot Domain



Crowd-sourced Annotations

- Insert pin in object
- Reach item in left drawer.
 - close the drawer
- Rub drawer handle
 - Do nothing
- open the top left drawer
- Reach to the dark green marker.
 - pull the drawer out
 - touch the tabletop
 - touch surface of table
- Hit the rack with the marker



Experiment 3: Is LOReL effective on a real robot?

“Open the left drawer”

Task (10 Trials Each)	LOReL
“Open the left drawer”	90%



Experiment 3: Is LOReL effective on a real robot?

“Open the right drawer”

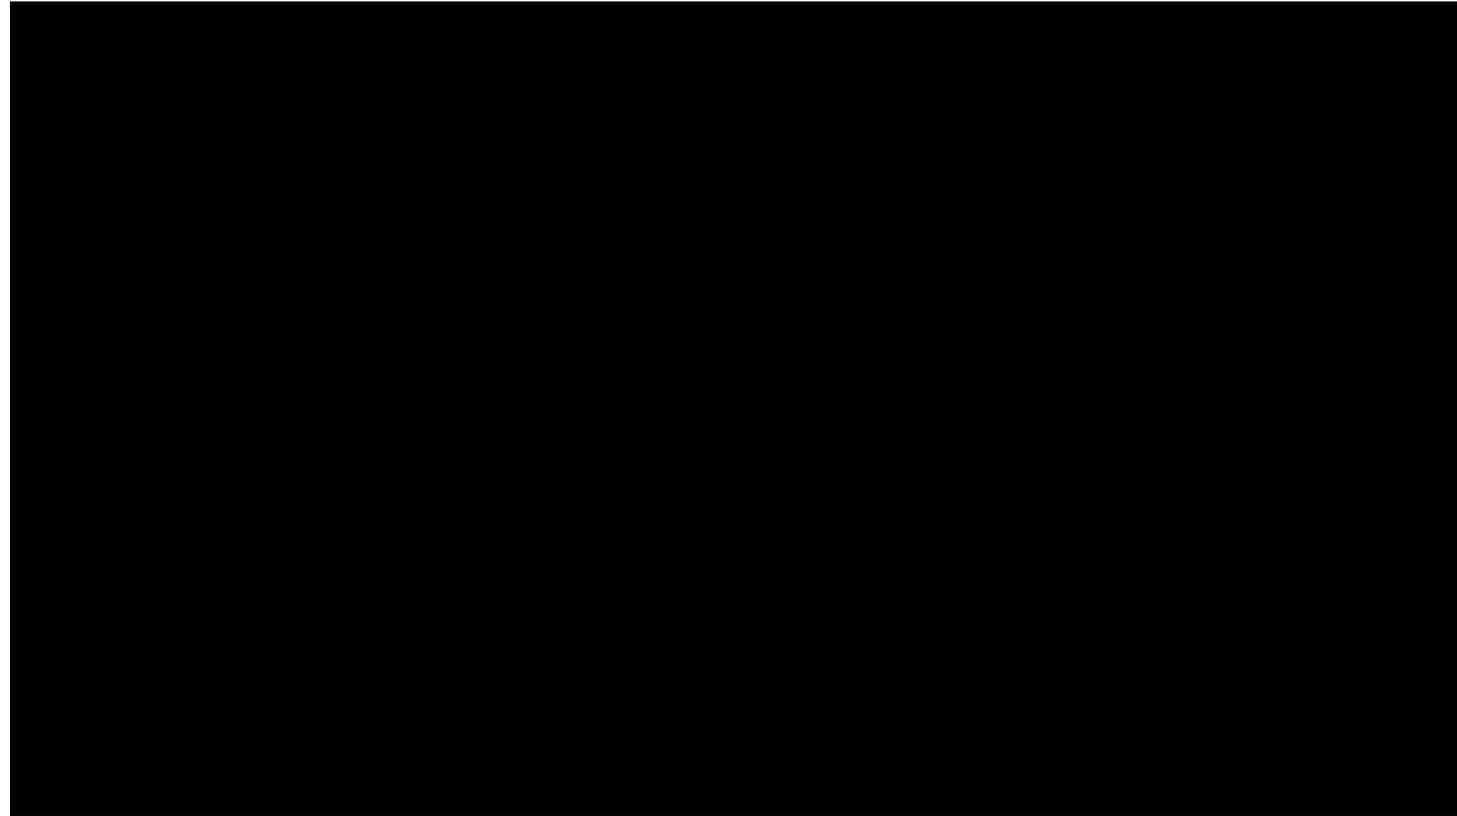


Task (10 Trials Each)	LOReL
“Open the left drawer”	90%
“Open the right drawer”	40%

Experiment 3: Is LOReL effective on a real robot?

“Move the stapler”

Task (10 Trials Each)	LOReL
“Open the left drawer”	90%
“Open the right drawer”	40%
“Move the stapler”	50%



Experiment 3: Is LOReL effective on a real robot?

“Reach the marker”

Task (10 Trials Each)	LOReL
“Open the left drawer”	90%
“Open the right drawer”	40%
“Move the stapler”	50%
“Reach the marker”	70%



Experiment 3: Is LOReL effective on a real robot?

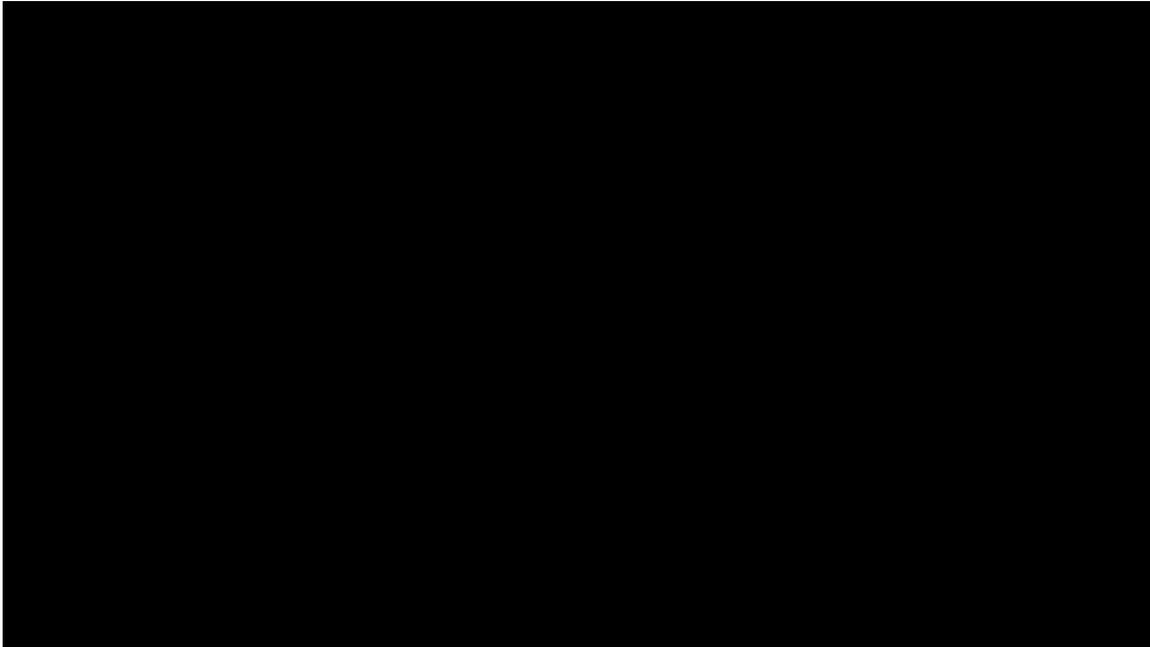
“Reach the cabinet”



Task (10 Trials Each)	LOReL
“Open the left drawer”	90%
“Open the right drawer”	40%
“Move the stapler”	50%
“Reach the marker”	70%
“Reach the cabinet”	80%
Average over tasks	66%

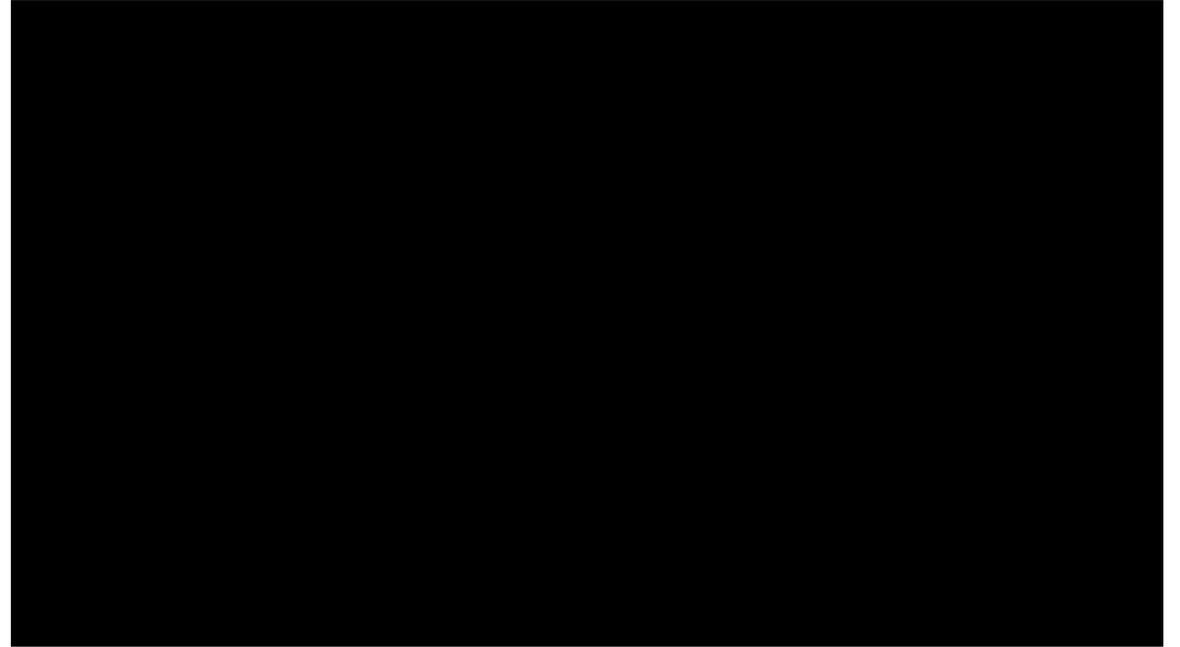
Robustness to Instruction Complexity

“Open the small black and white drawer on the left fully”



70% Success Rate

“Push the small gray stapler around on top of the black desk”



50% Success Rate

Key Takeaways

To learn useful multi-task robot policies humans need effective task-specification like natural language

Key Takeaways

To learn useful multi-task robot policies humans need effective task-specification like natural language

By annotating sub-optimal offline data with crowd-sourced instructions, LOReL can learn language-conditioned visuomotor skills on a real robot

Key Takeaways

To learn useful multi-task robot policies humans need effective task-specification like natural language

By annotating sub-optimal offline data with crowd-sourced instructions, LOReL can learn language-conditioned visuomotor skills on a real robot

By using pre-trained language models LOReL can generalize to unseen natural language instructions + handle crowdsourced annotations