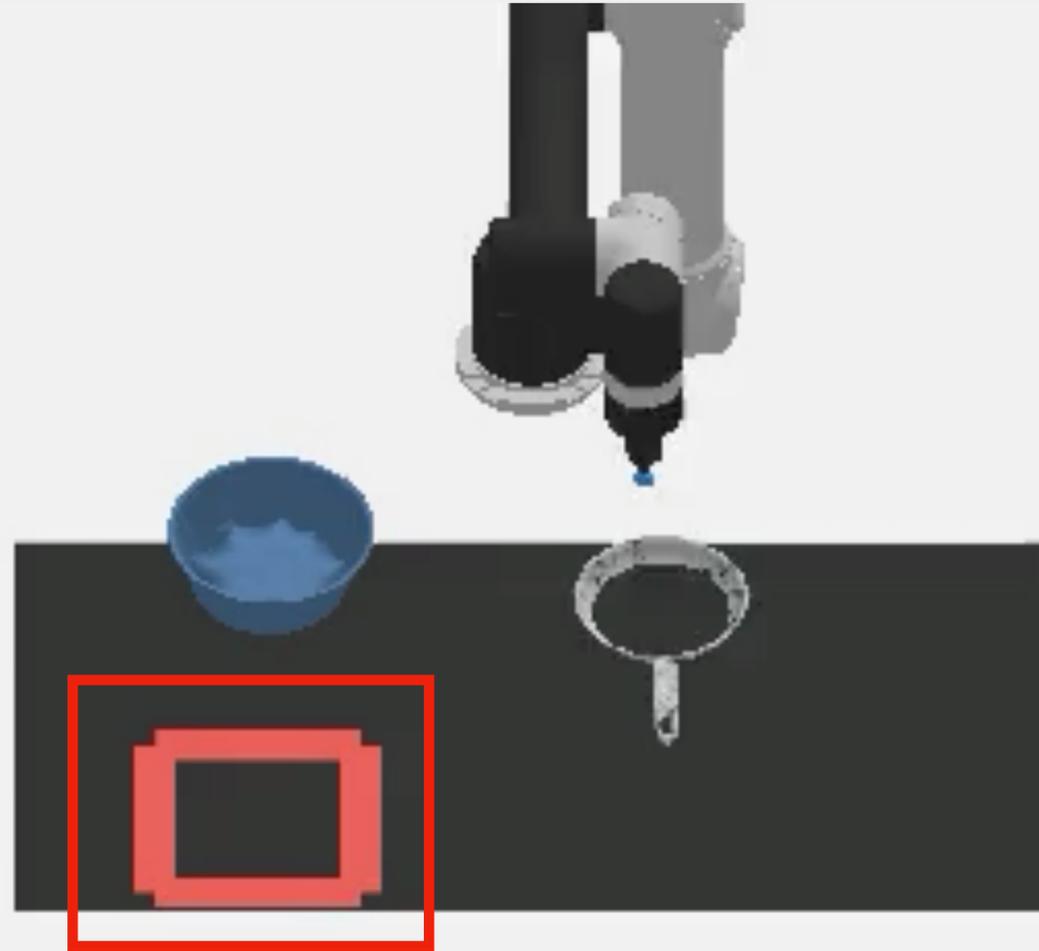


Learning with Language-Guided State Abstractions

Andi Peng, Ilia Sucholutsky*, Belinda Li*, Theodore Sumers,
Thomas Griffiths, Jacob Andreas*, Julie Shah*





Grabbing a pan?

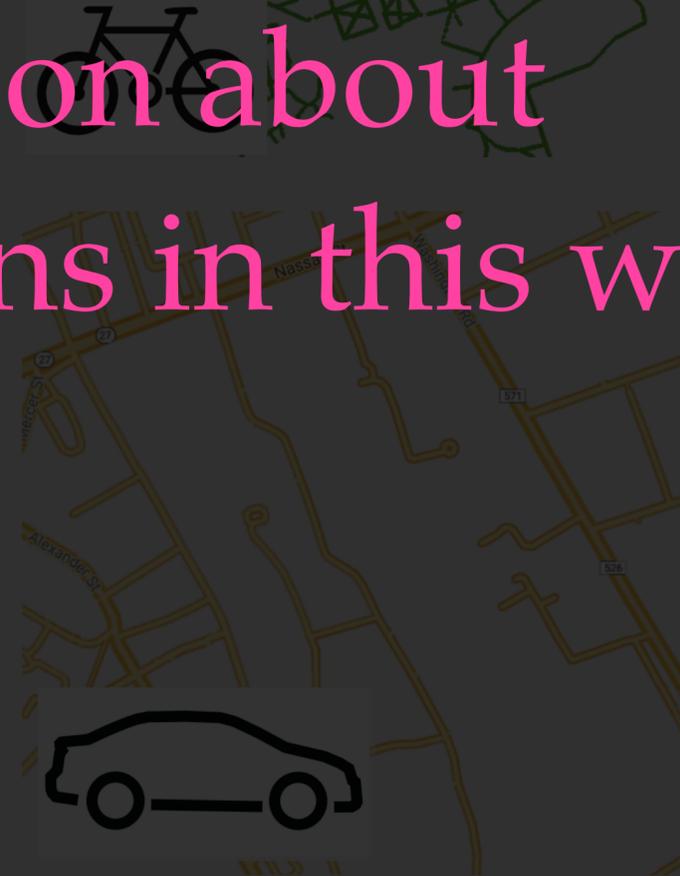
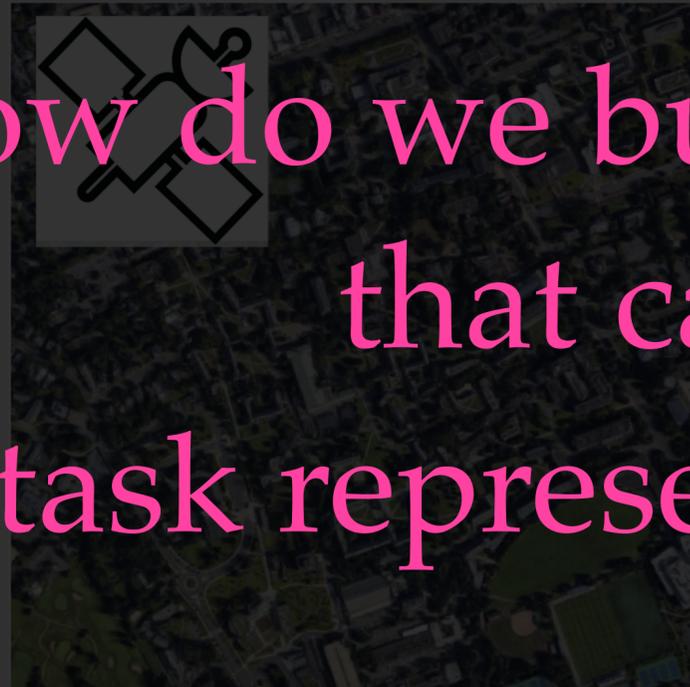
Grabbing a metal object?

Putting any object
on the stove?

Cooking

Humans plan using **state abstractions**

How do we build autonomous agents
that can reason about
task representations in this way?

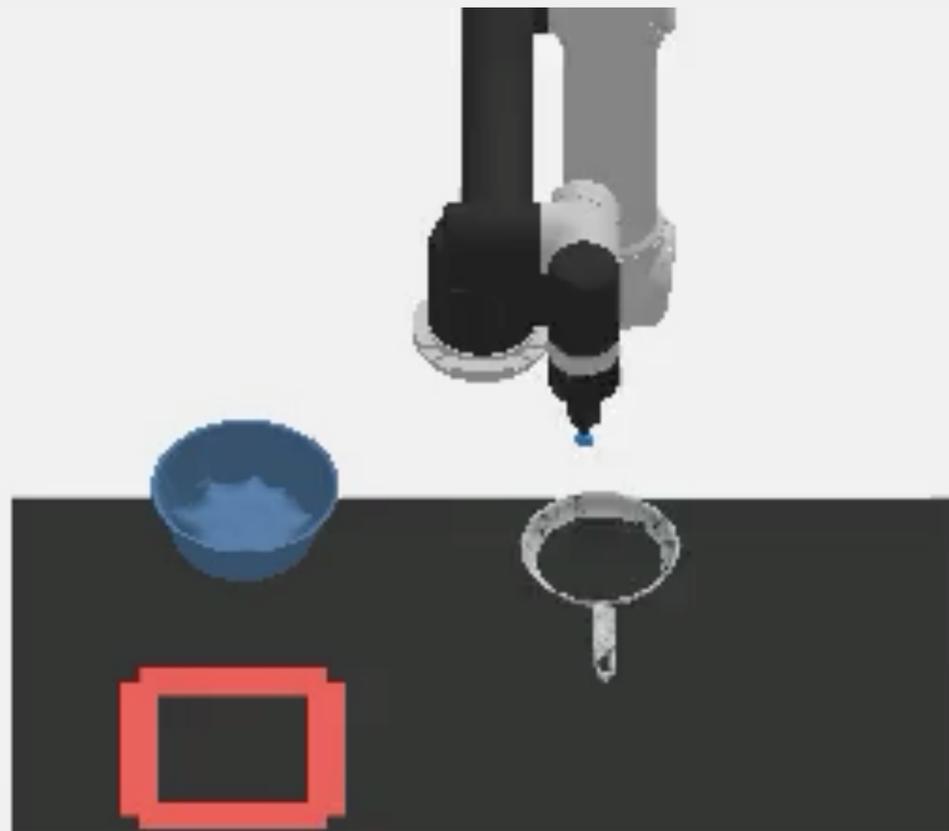


Key Insight

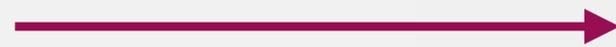
Language serves as a natural **semantic prior**
to help construct *abstractions*.

How do we do this?

Idea: Query language models for task-relevant features



Pick up the pan

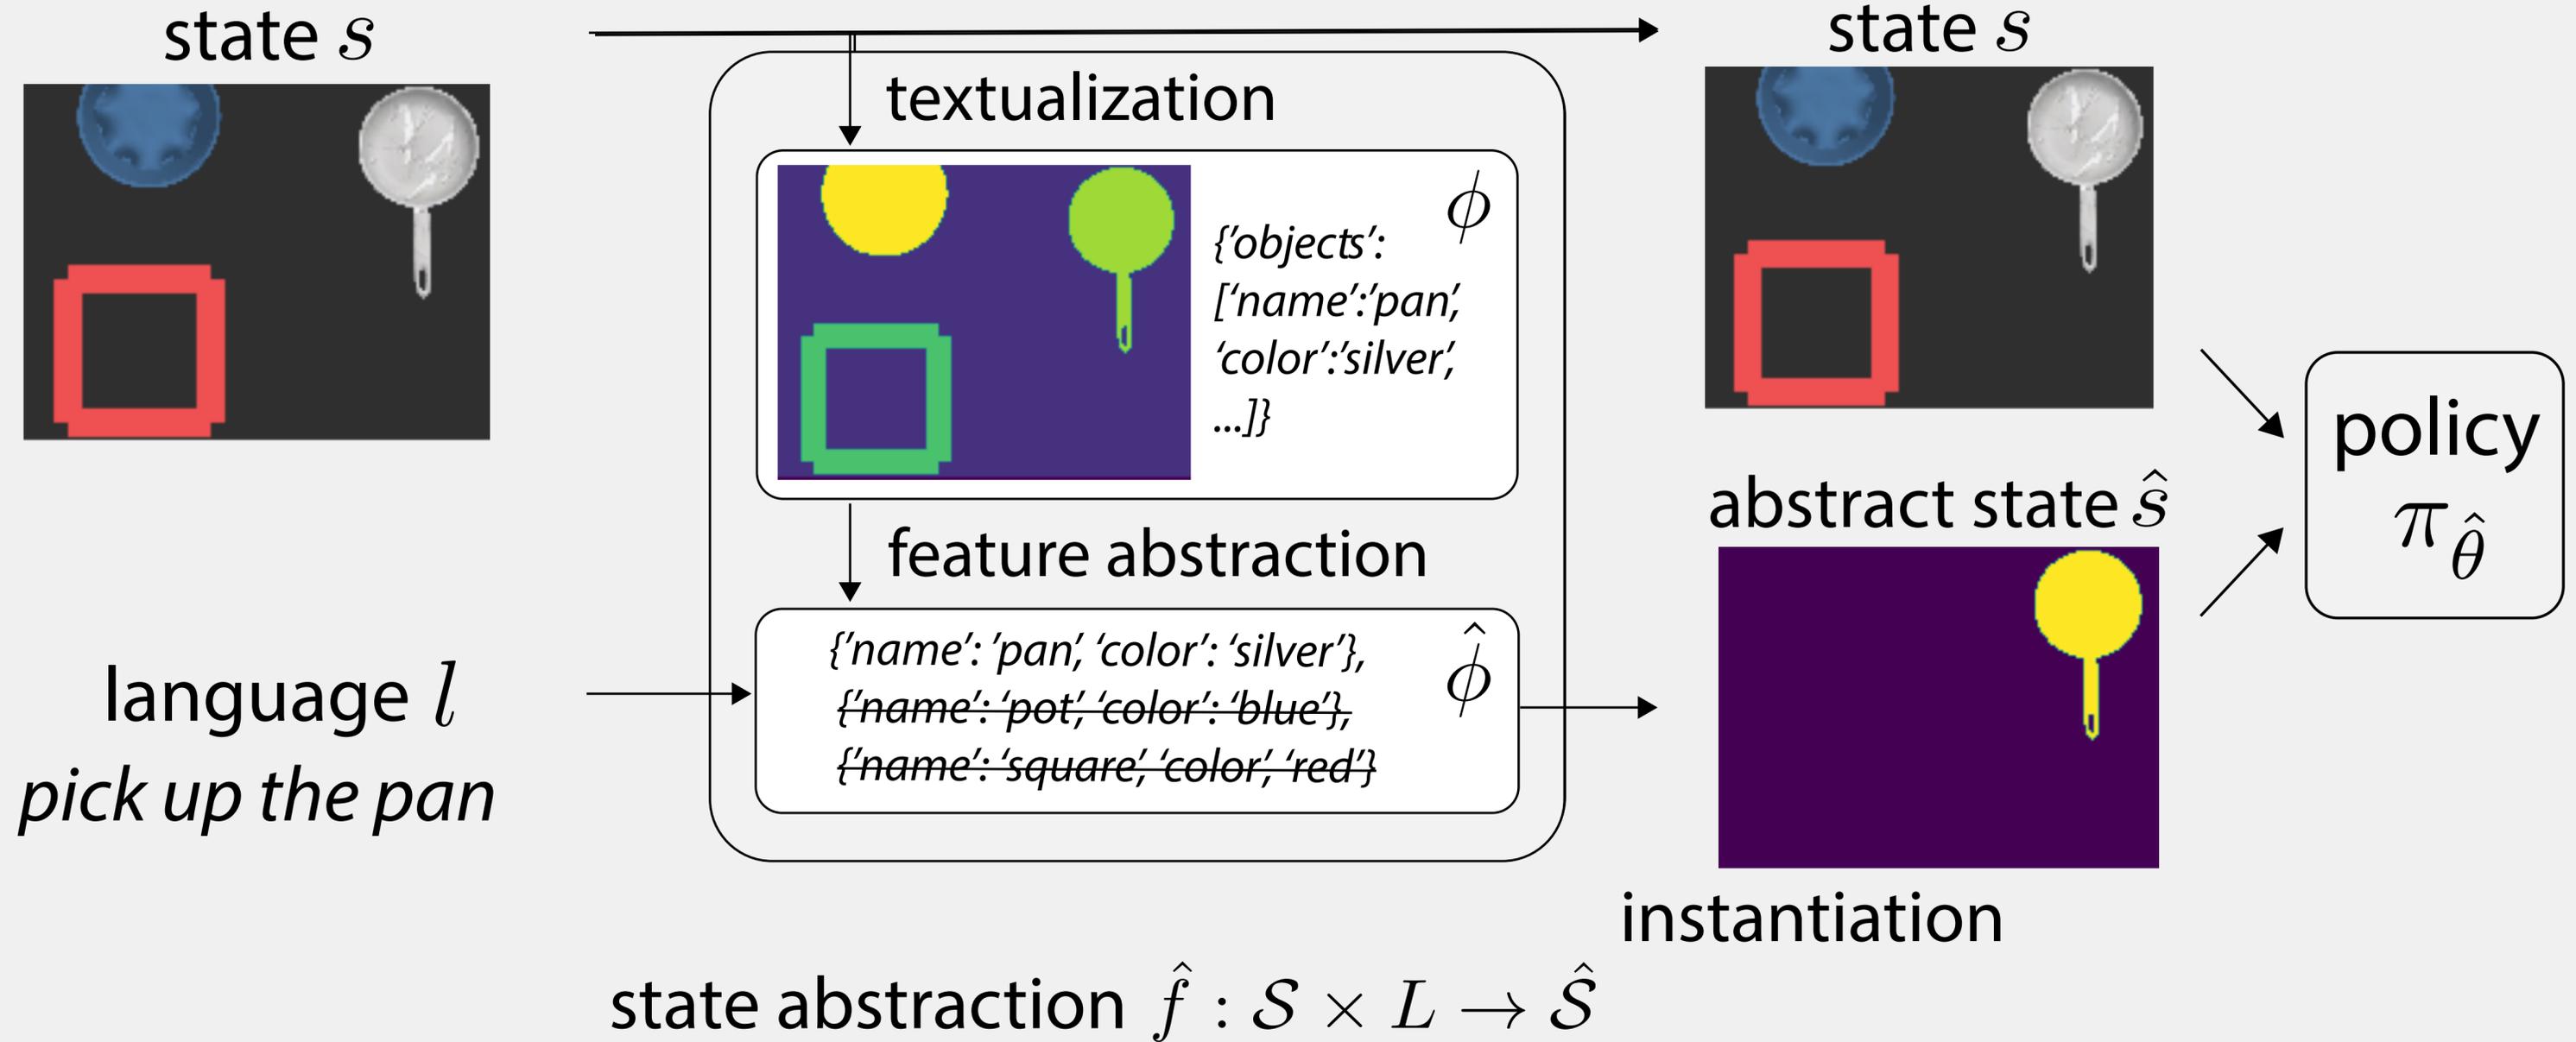


Q: Is *color* of the pan important to the task?

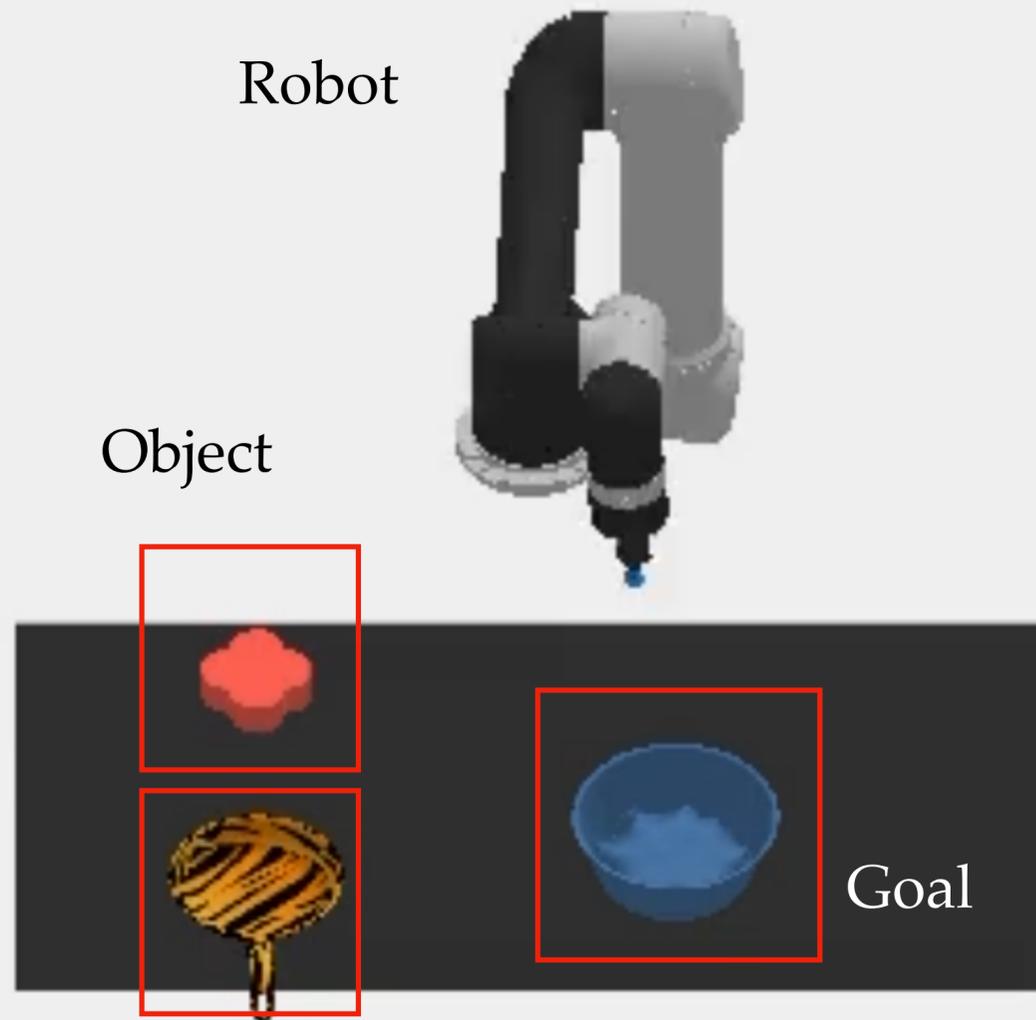
Q: Is the object being a *pan* important to the task?

Q: Is the *bowl* important to the task?

LGA: Language-Guided Abstraction



We evaluate on different tasks in VIMA



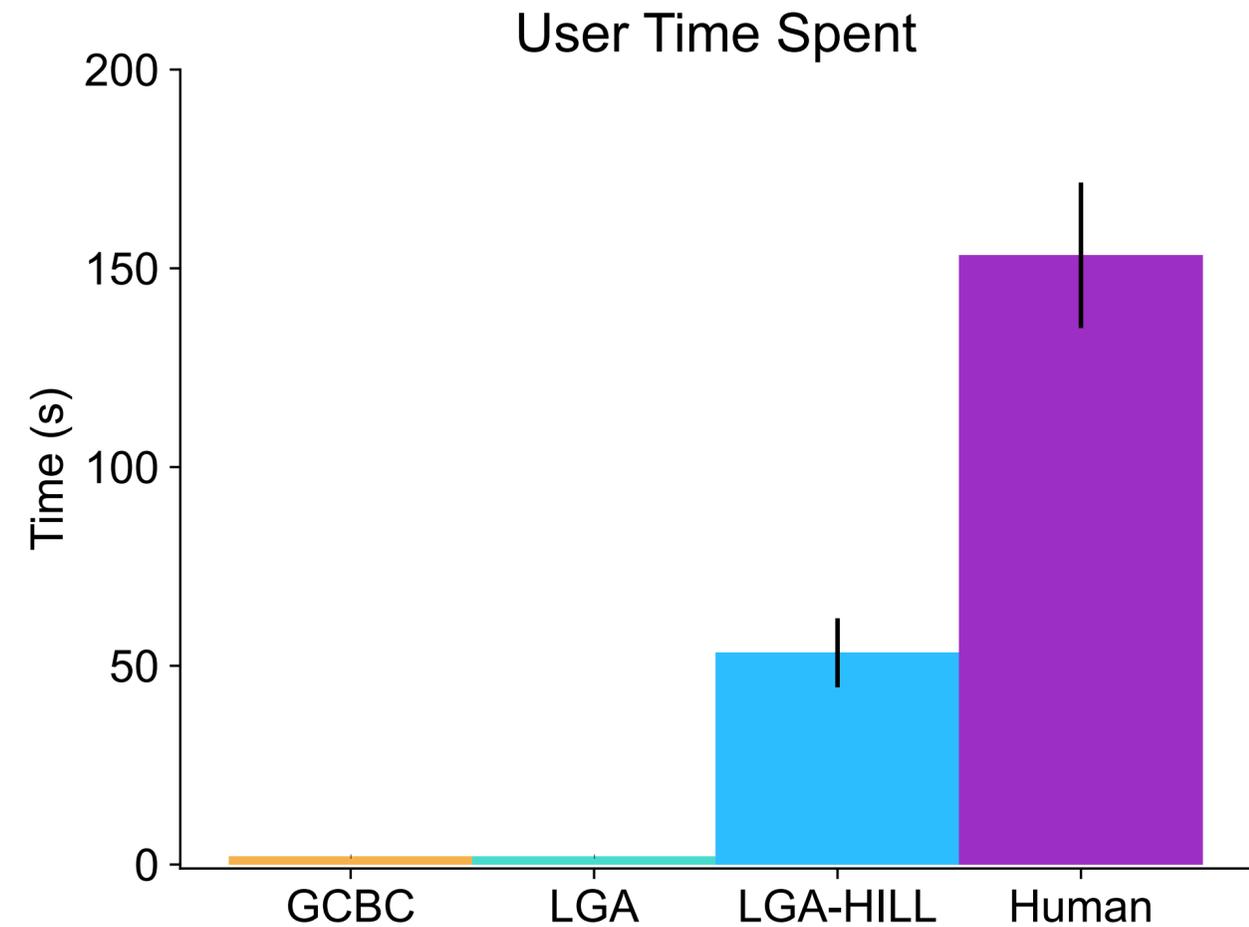
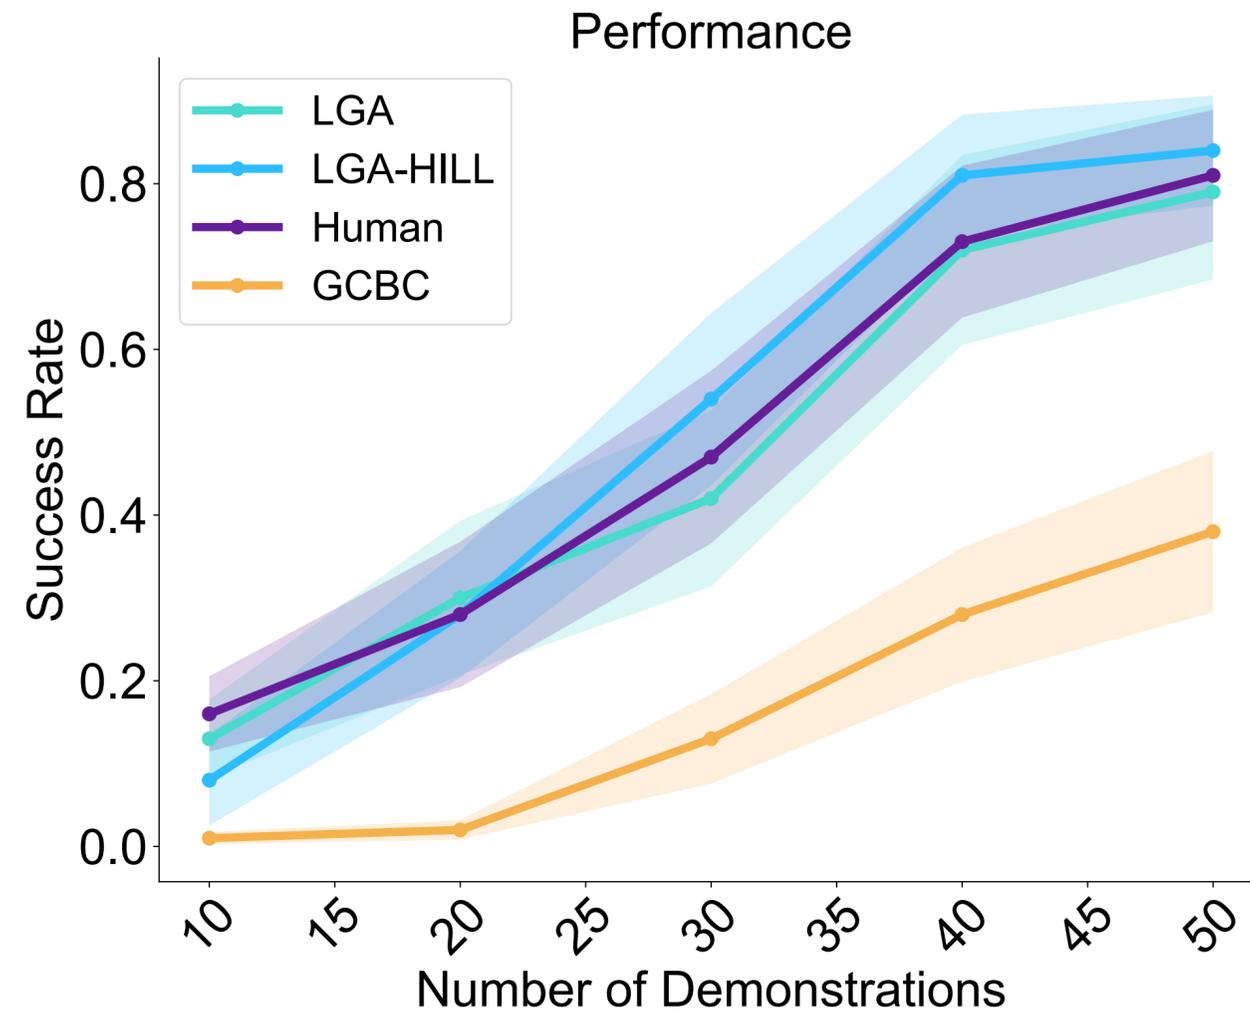
Pick up the pan

Pick up the the red pan

Pick up something to drink from

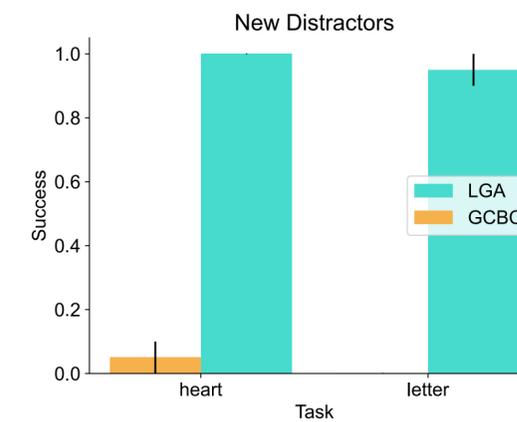
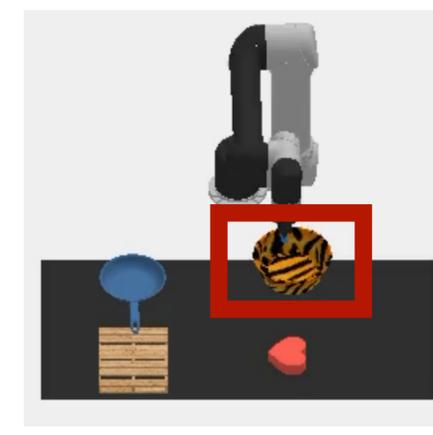
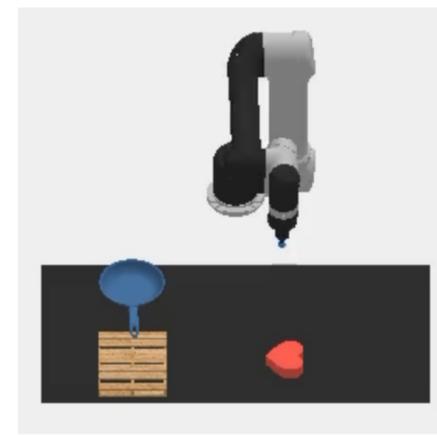
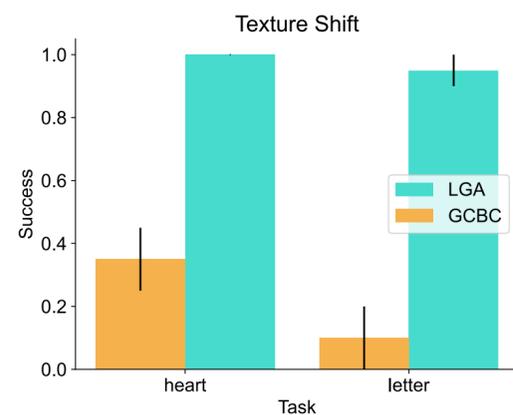
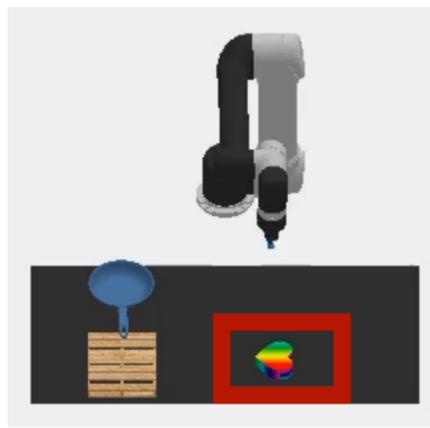
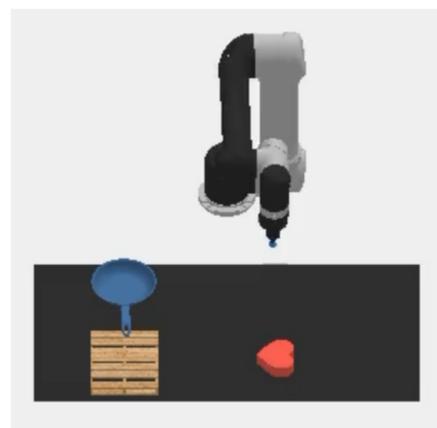
Possible Distractor

Does LGA construct good state abstractions?



LGA outperforms naive BC on task performance while significantly reducing user time spent compared to manual feature specification ($p < 0.001$).

Does LGA improve policy robustness?

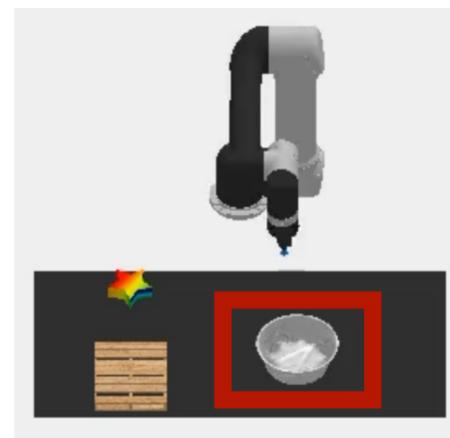
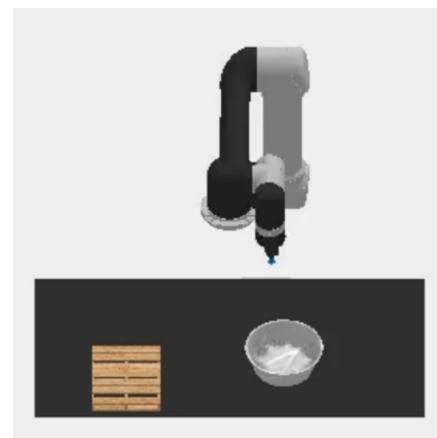


Bring me the heart.

Bring me the heart.

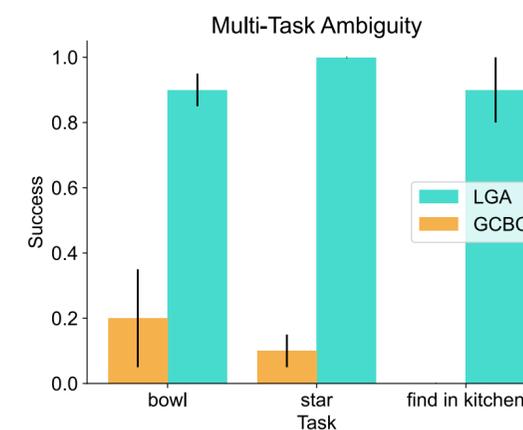
Bring me the heart.

Bring me the heart.



Bring me the bowl.

Bring me something I can find in a typical kitchen.



LGA is more robust to 1) state covariate shifts and 2) linguistic ambiguity.

Learning with Language-Guided State Abstractions

Andi Peng, Ilia Sucholutsky*, Belinda Li*, Theodore Sumers,
Thomas Griffiths, Jacob Andreas*, Julie Shah*

