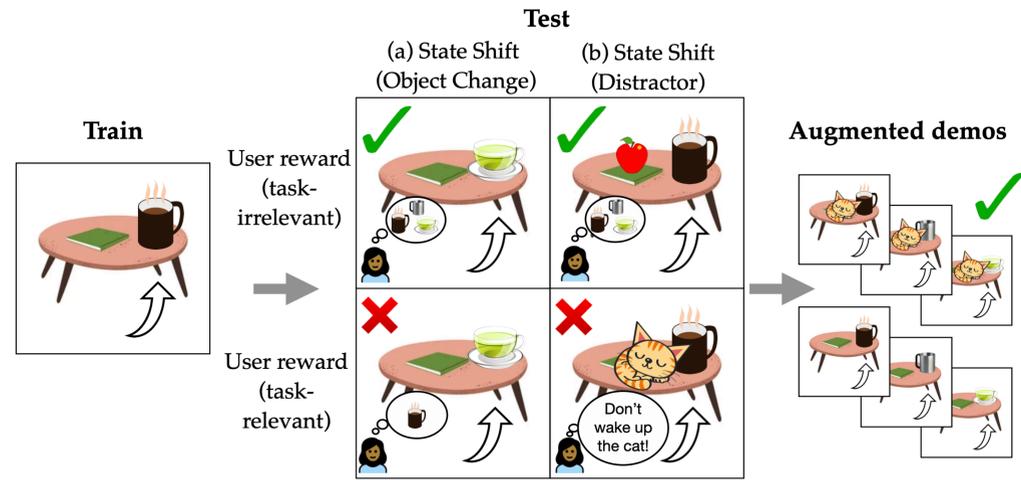


Diagnosis, Feedback, Adaptation:

A Human-in-the-Loop Framework for Test-Time Policy Adaptation

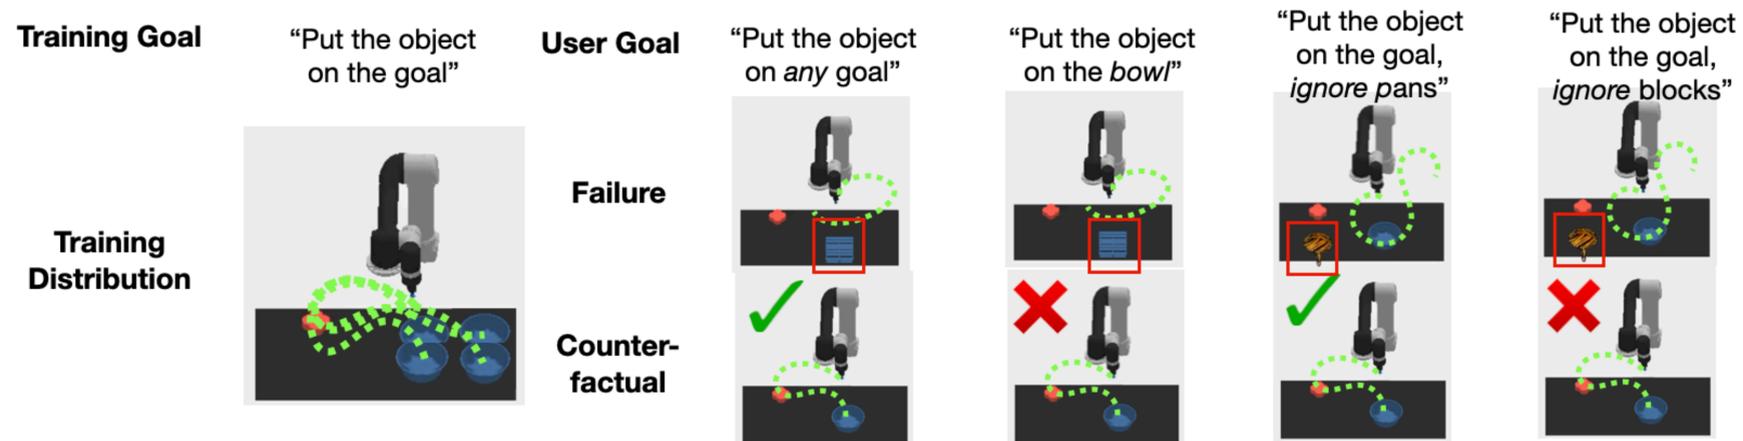
Andi Peng, Aviv Netanyahu, Mark Ho, Tianmin Shu, Andreea Bobu, Julie Shah, Pulkit Agrawal

Problem: Data augmentation can help robot policies be more robust to distribution shift by varying *task-irrelevant* concepts, but how do we know what is irrelevant?

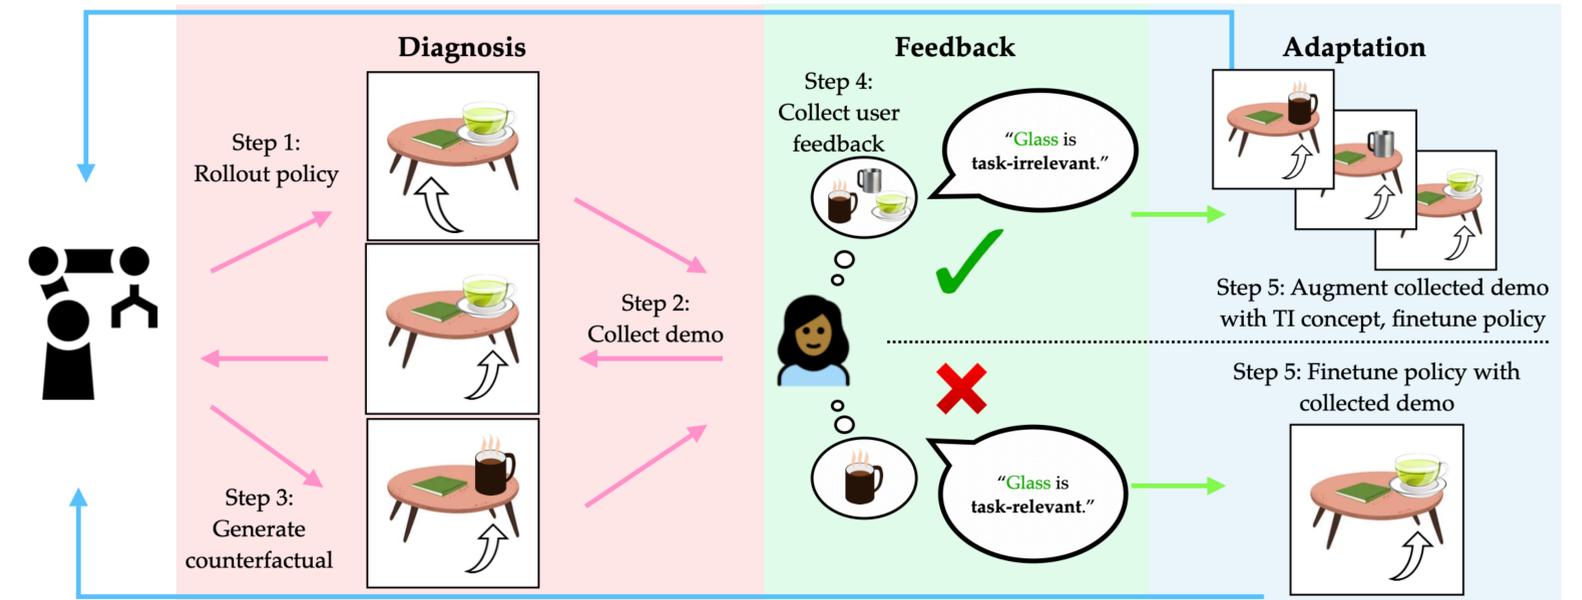


Key insight: End users are *uniquely positioned* to know what is *task-irrelevant*.

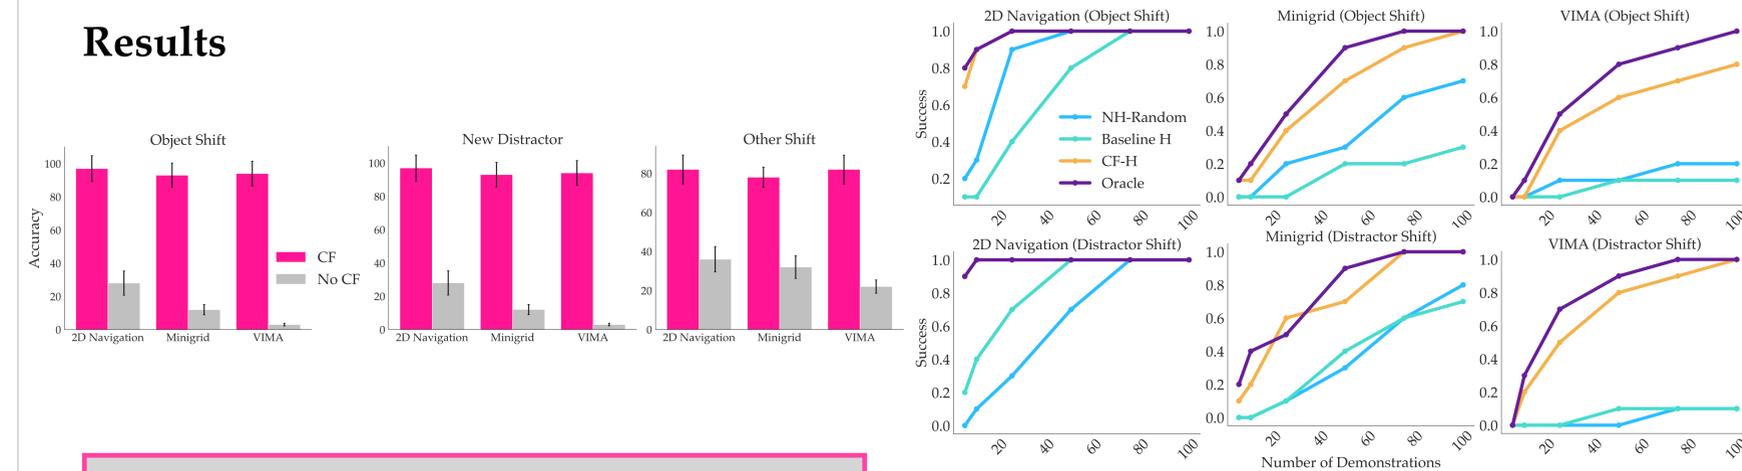
Idea: Provide *counterfactual* demonstrations as a contrastive explanation to help users isolate the cause of failure.



Framework



Results



Users are significantly more accurate at identifying the correct shifted concept with our method than with behaviour alone.

Our method is more sample efficient than human and random augmentation baselines.