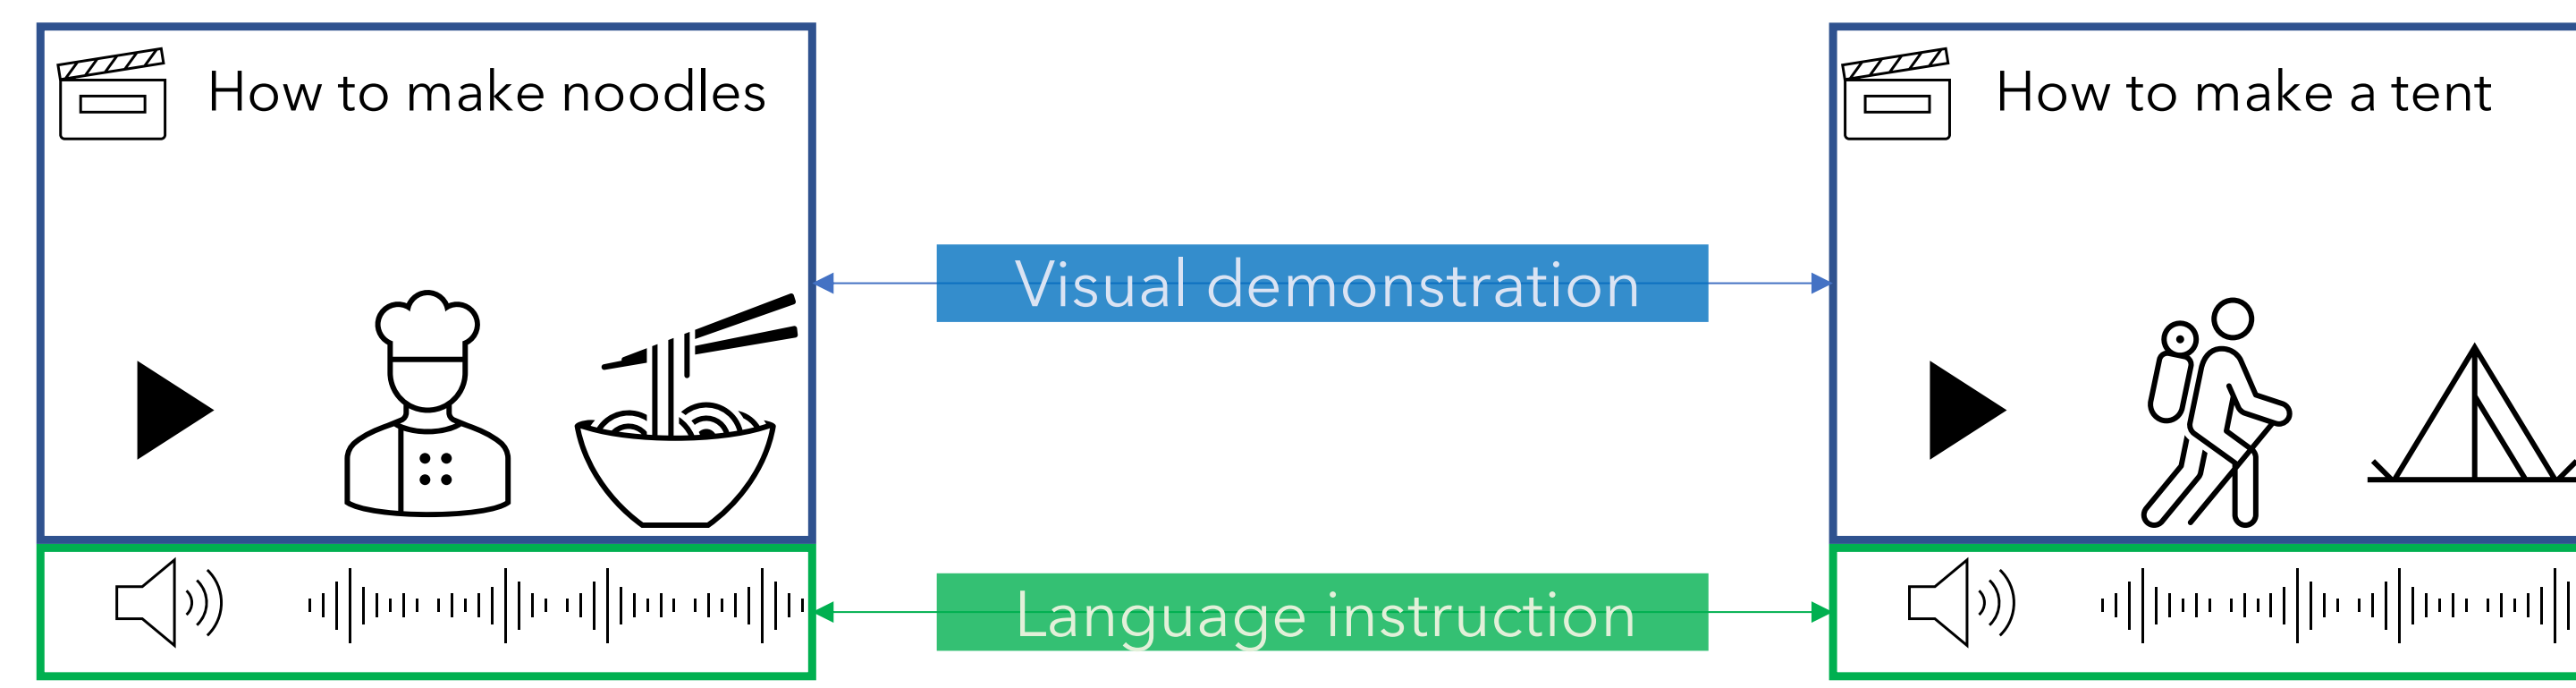



Why Condition Robotic Policies on Demonstrations + Language?

Learn Efficiently with Multimodality



Humans teach/learn with demos + language (e.g. instructional videos).

Resolve Ambiguities




Hard to infer a task from multiple demos. Instruction contains many ambiguities (*).


Specifying tasks with demos + language helps resolve ambiguities.

Language-only Task Specification

Instruction: "Grab the left* pointy* gold* object."



Overview

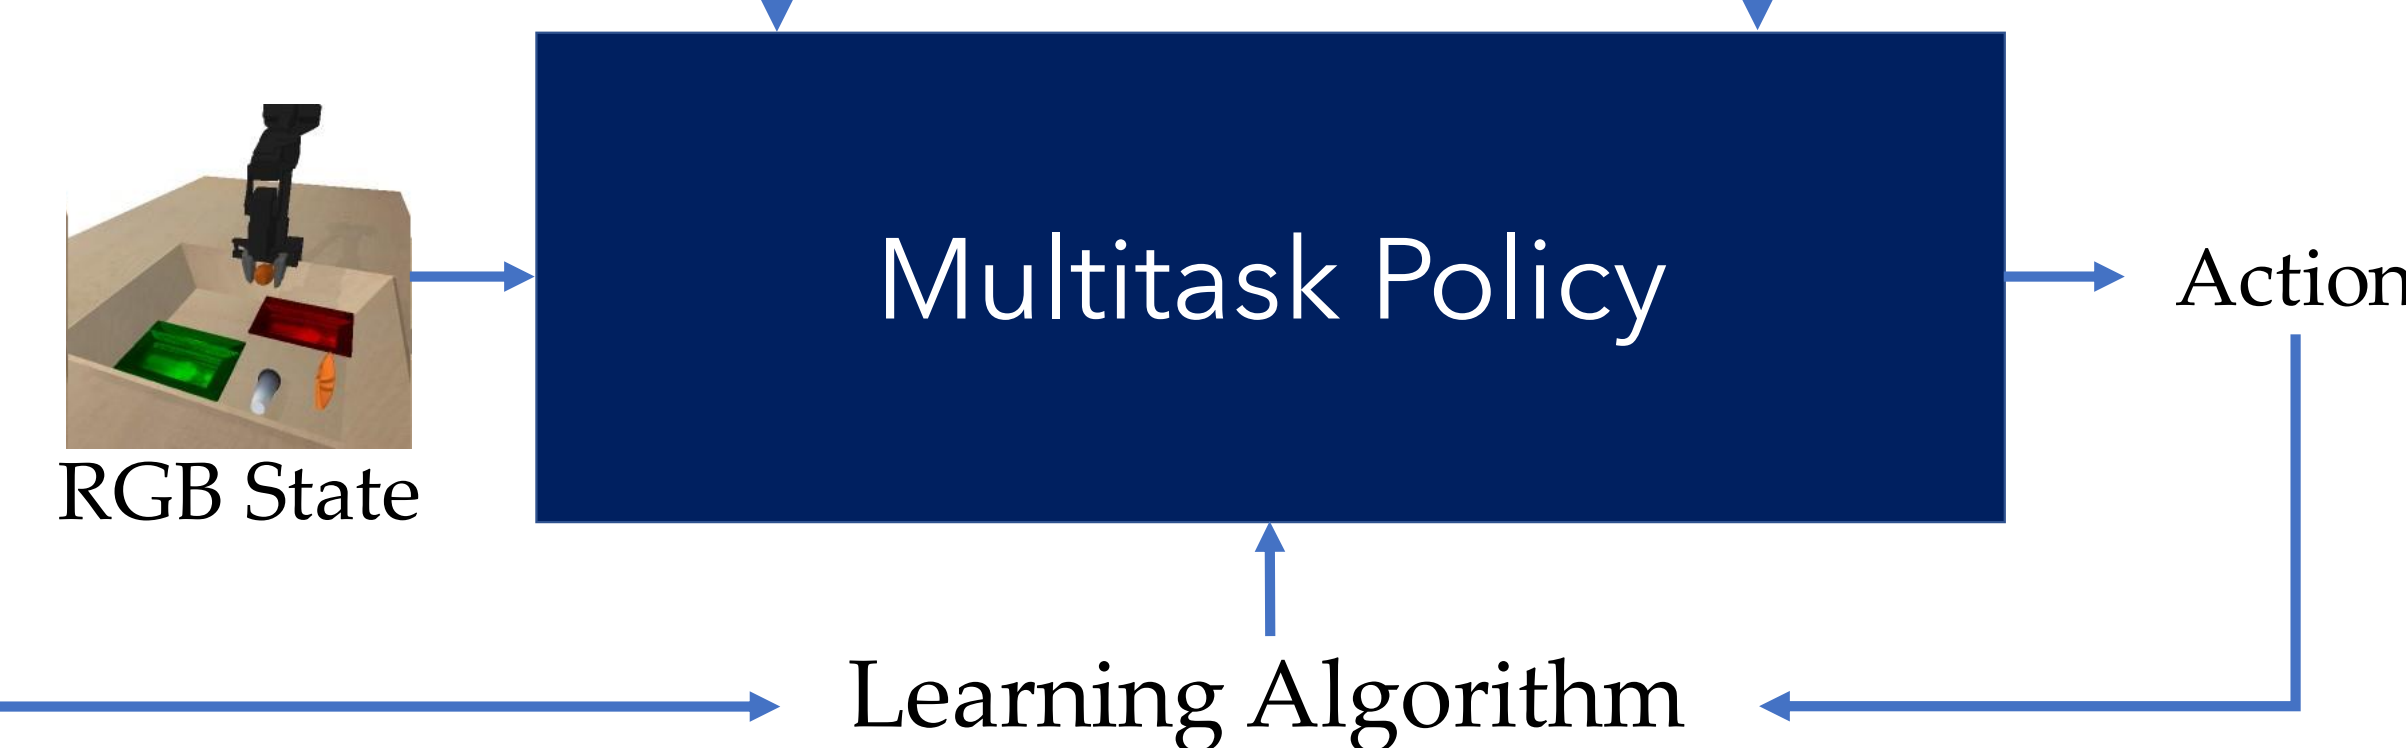
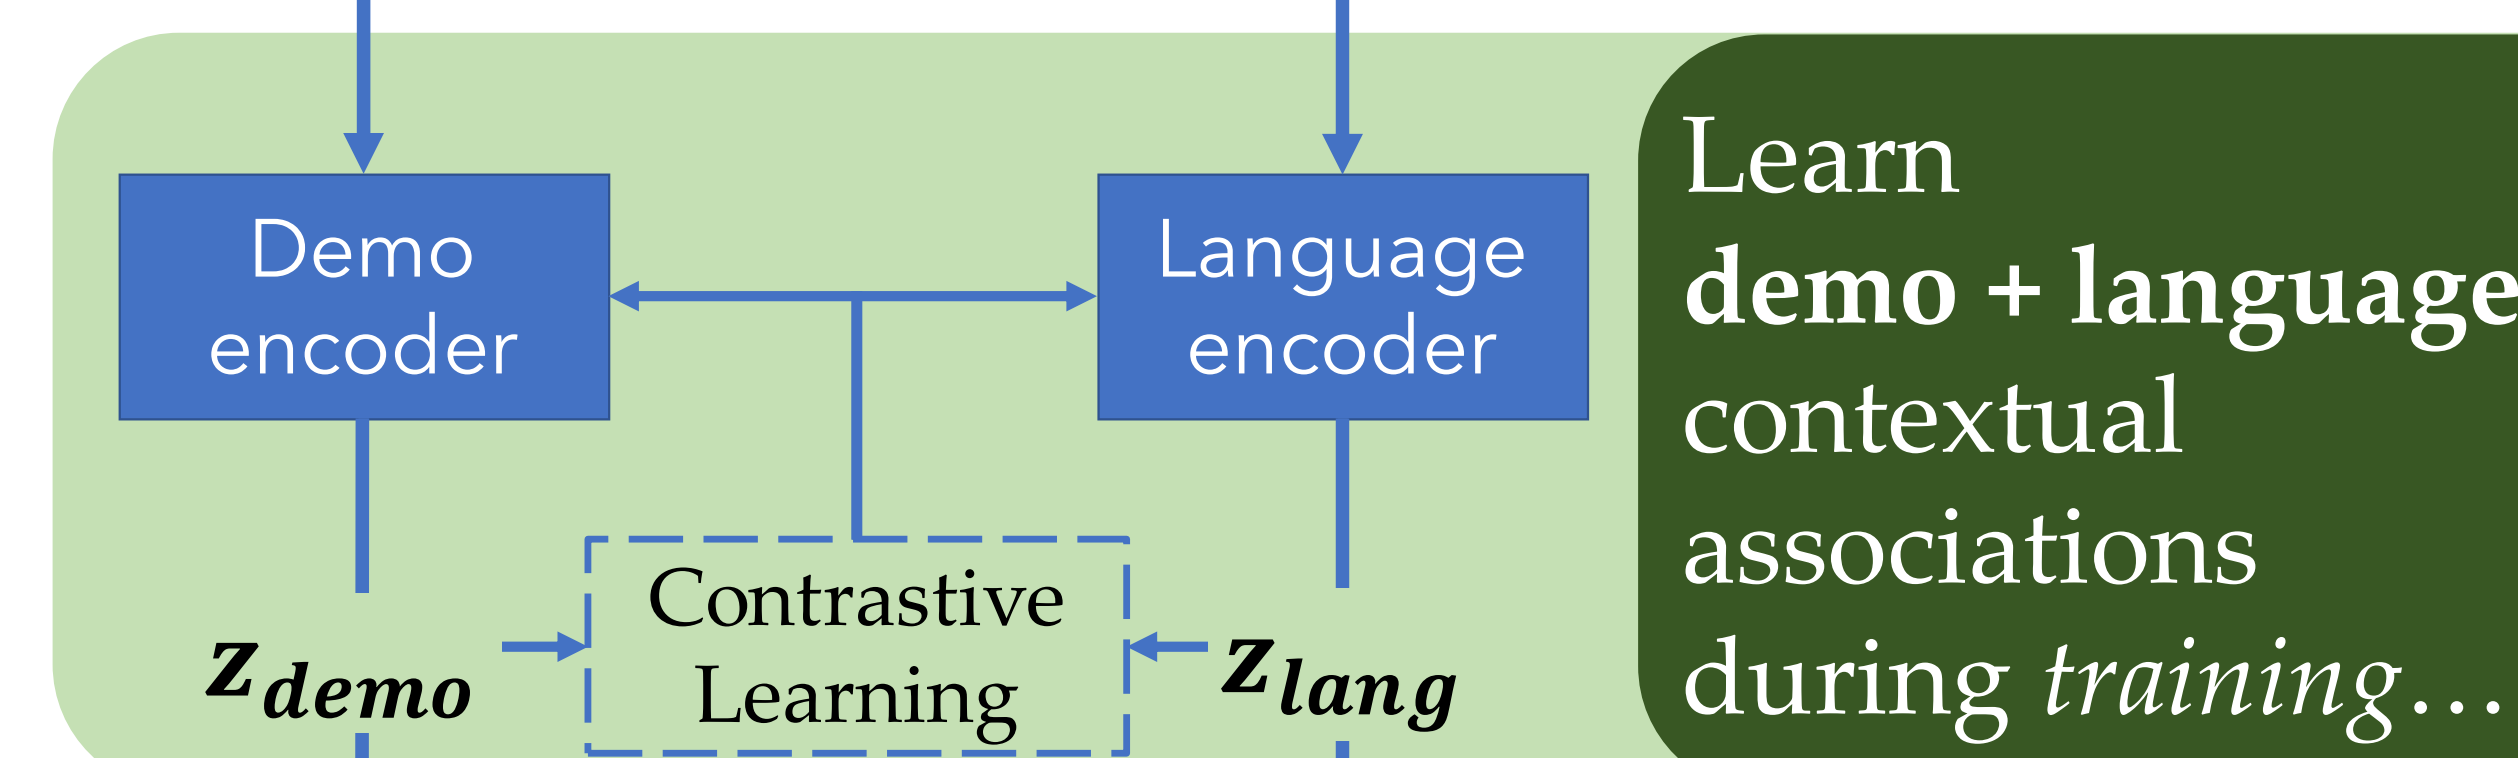
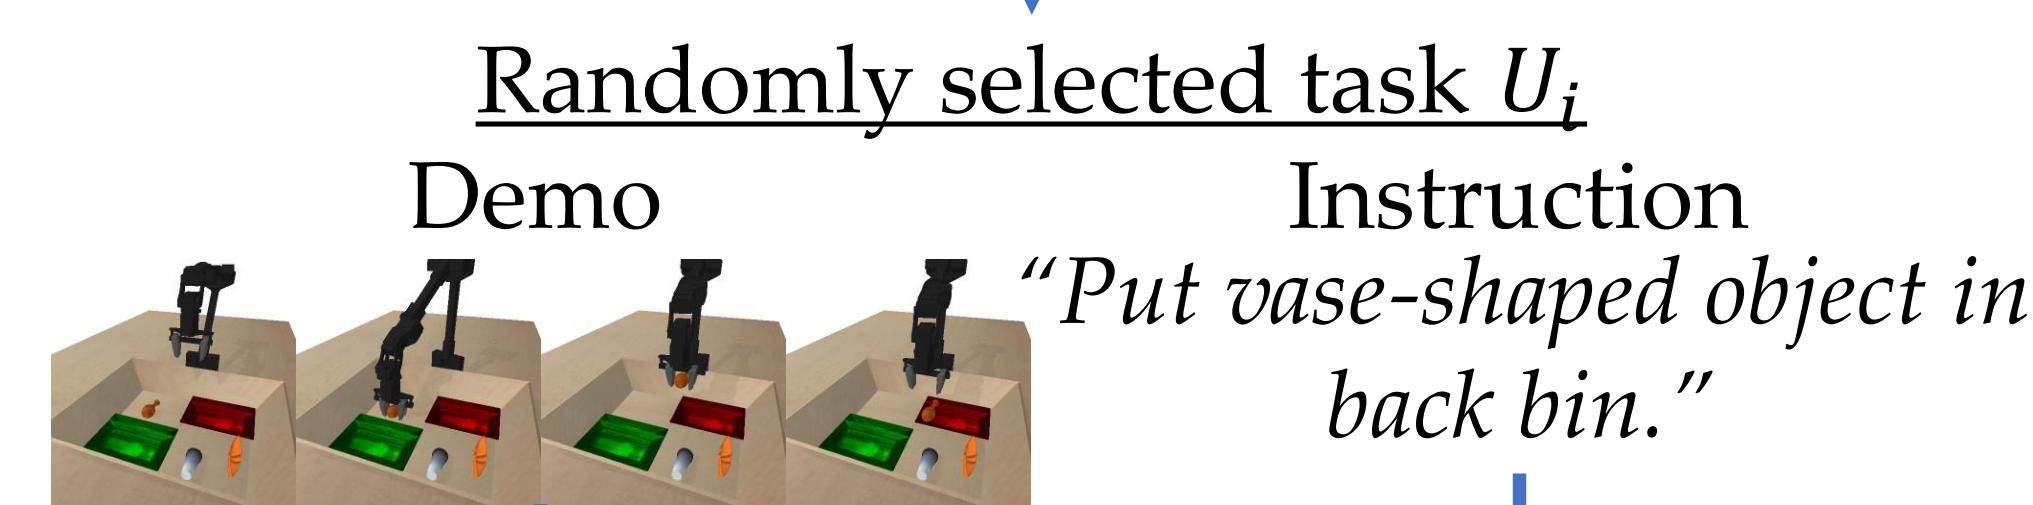
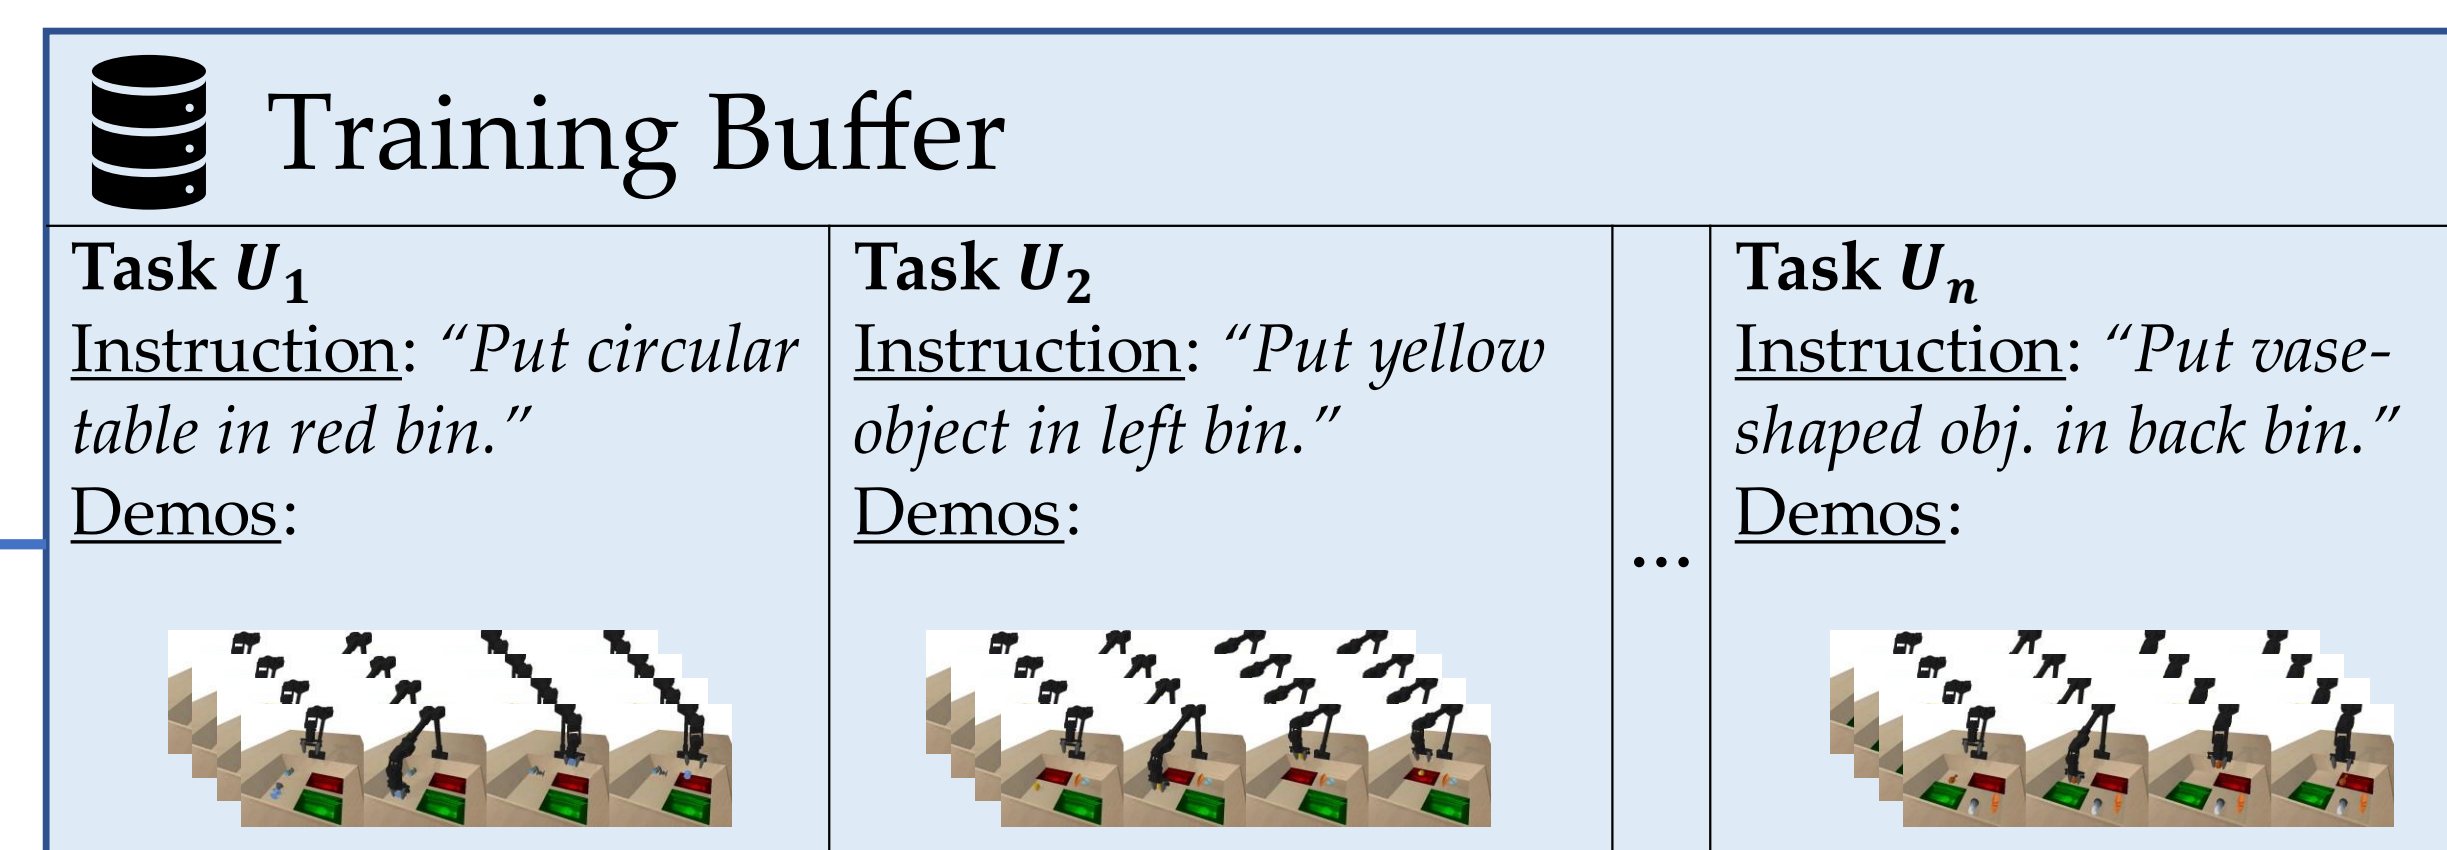


DeL-TaCo (Demo-Language Task Conditioning)

Simultaneously conditions on *both* demos *and* language to teach robots new tasks, unlike prior work.

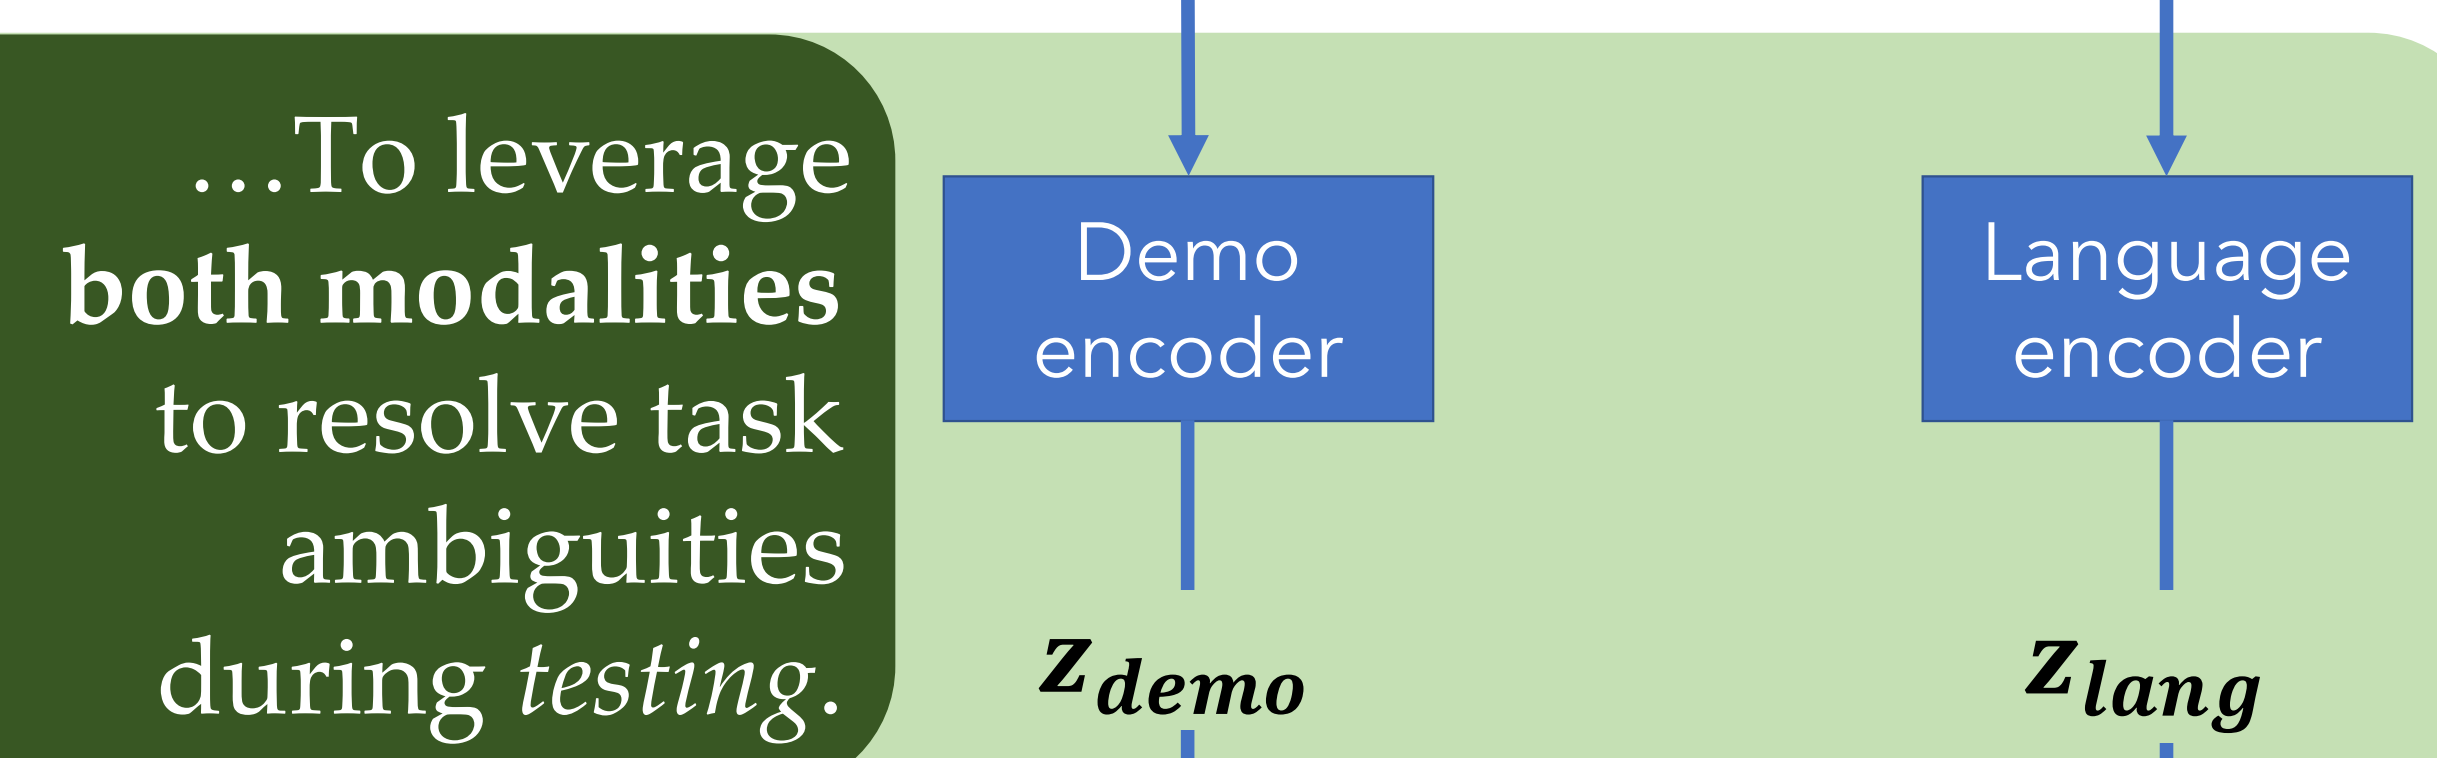
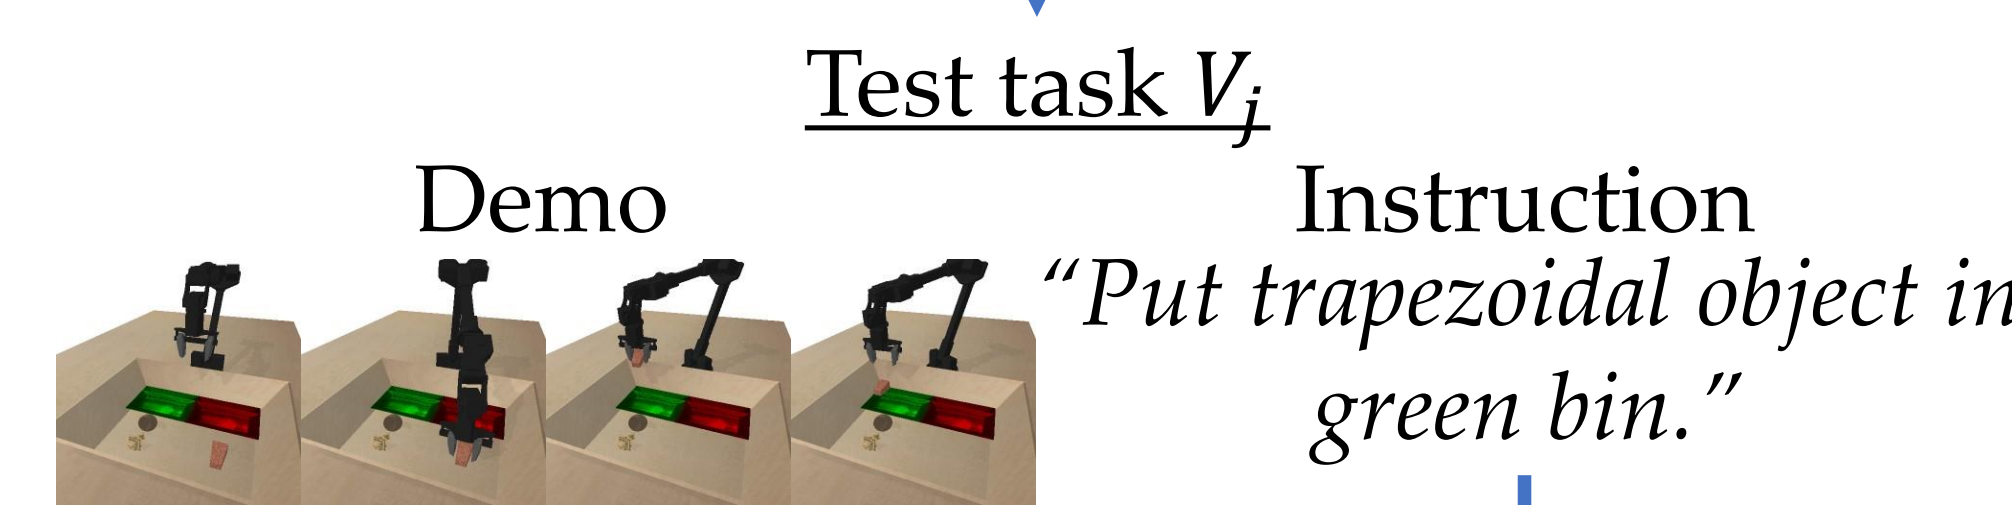
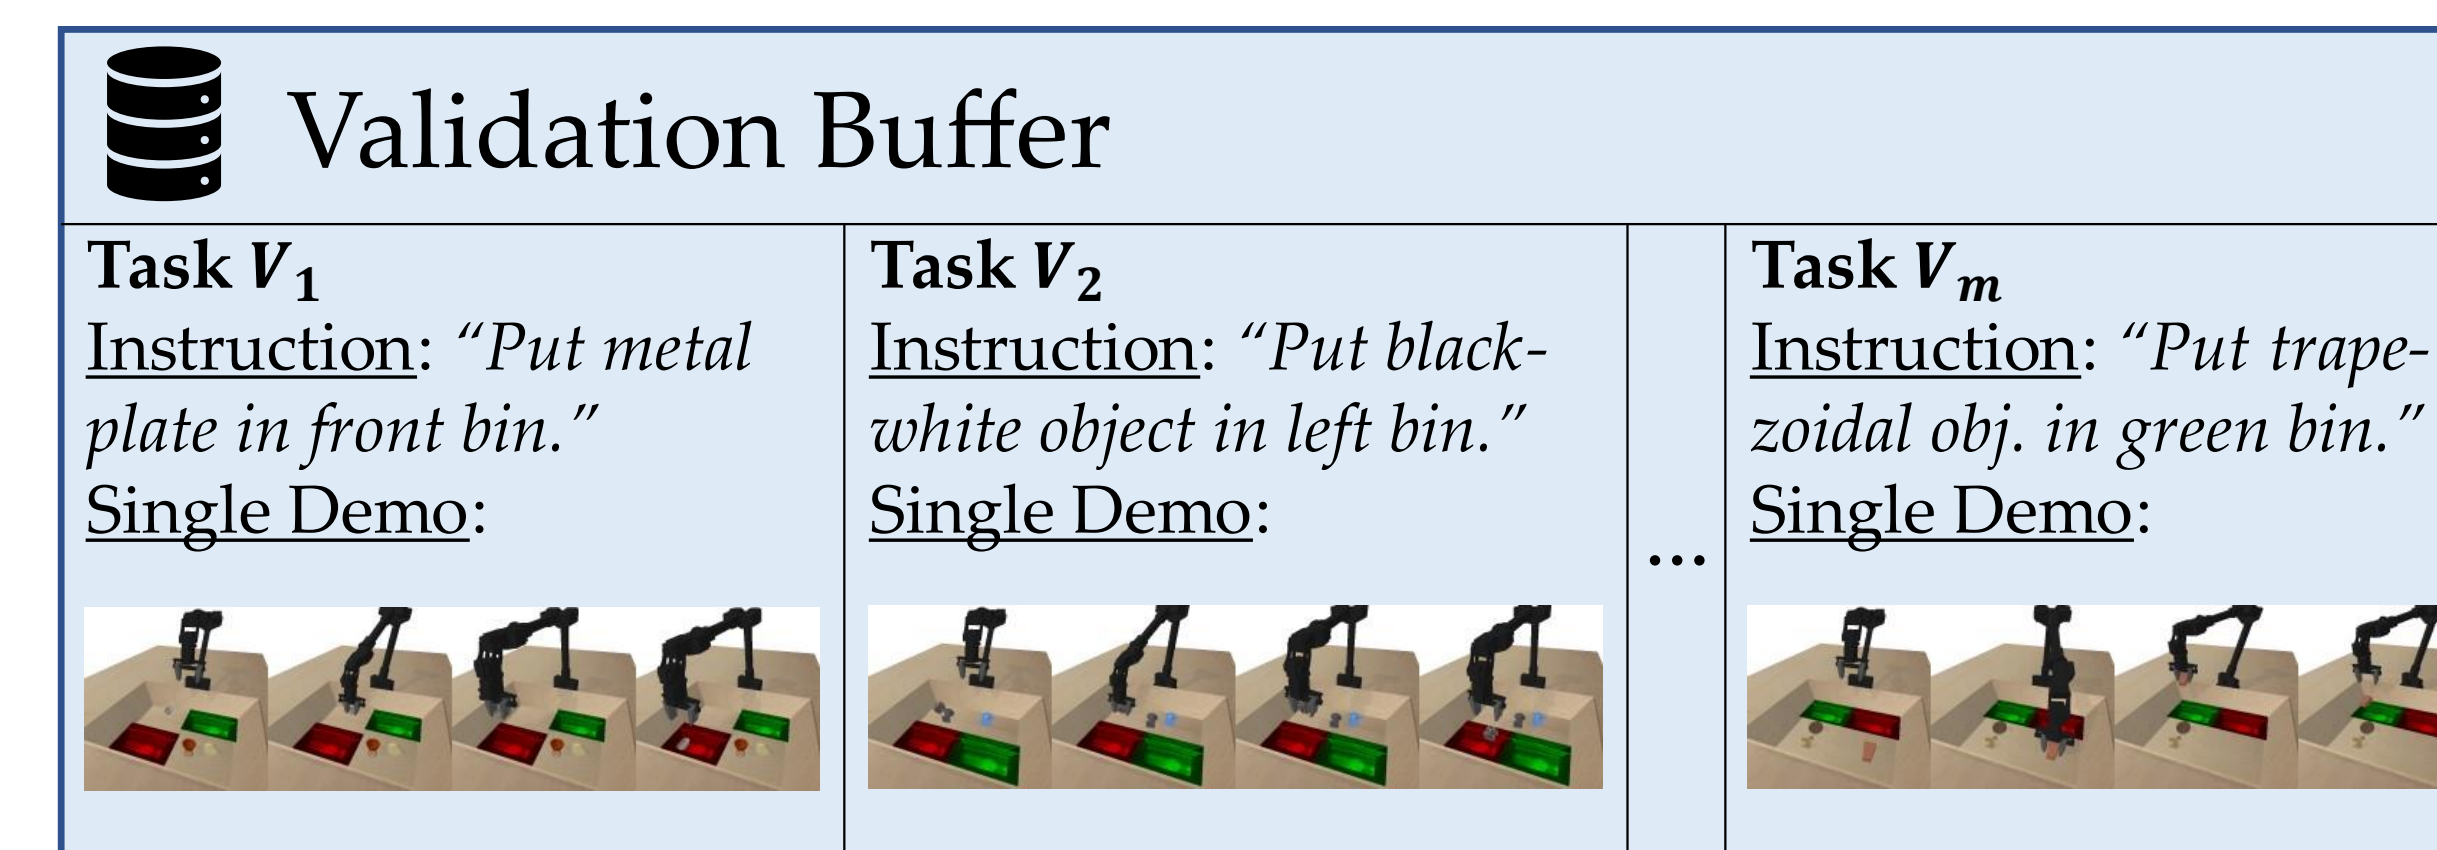
Training

Train a single multi-task policy on hundreds of tasks.



Testing

One-shot generalization to ~100 new tasks (new objects, colors, shapes).



Summary

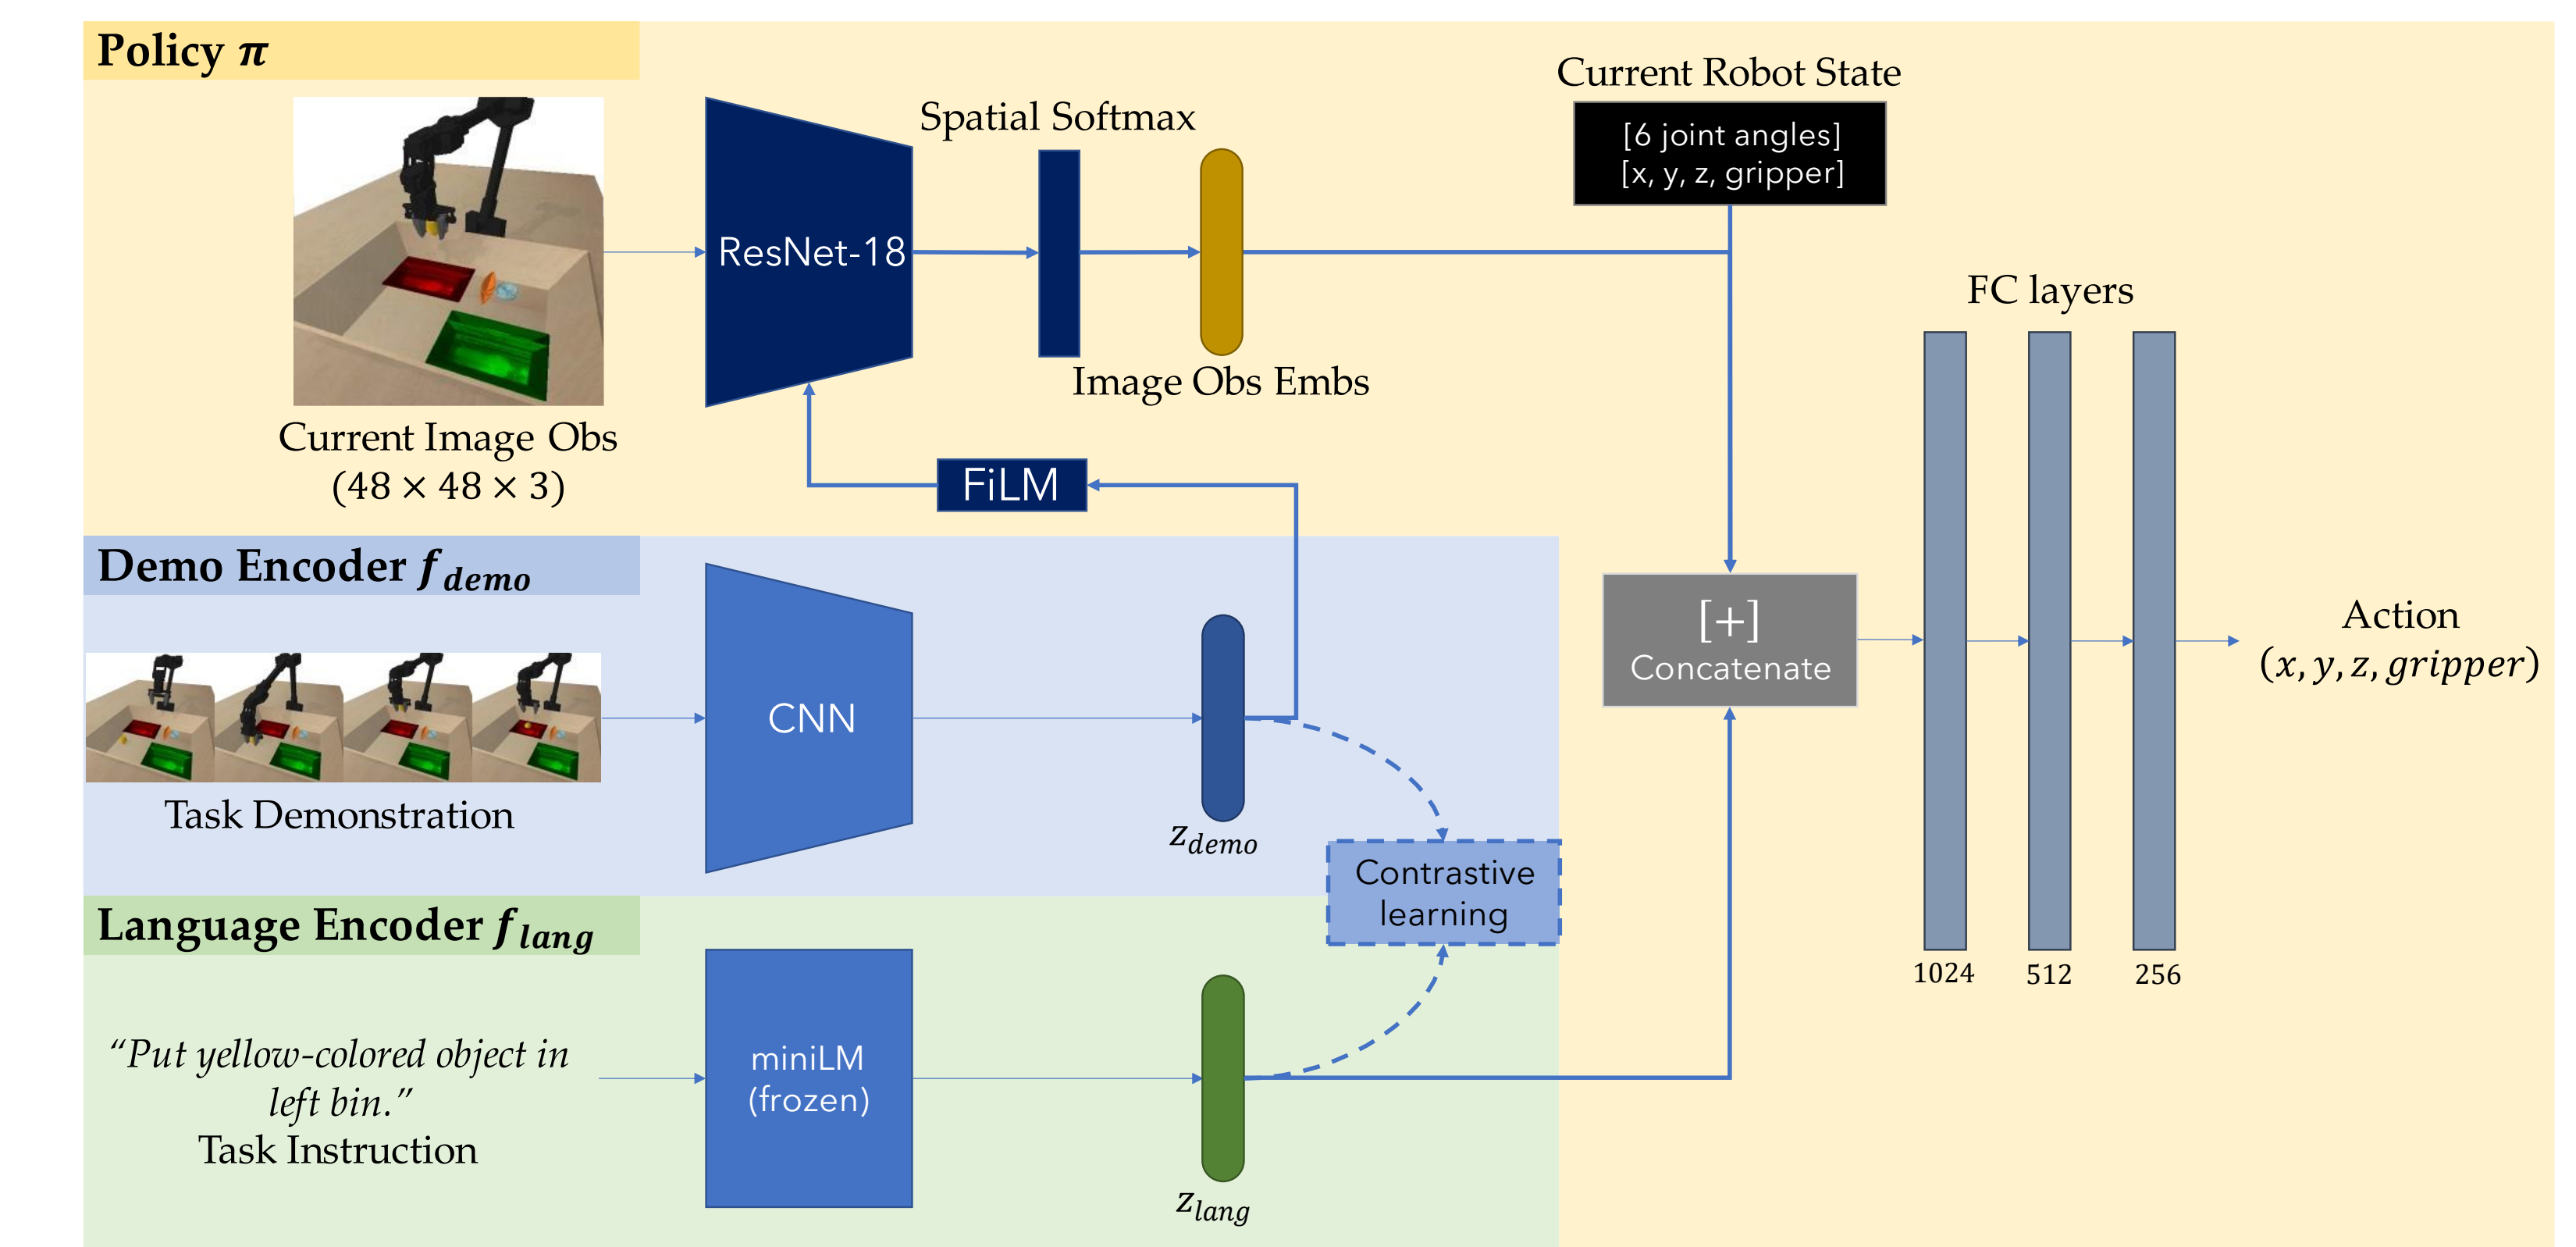
Contributions

- 1st work (we know of) showing that learning with demos + language attains better generalization and sample efficiency than prior (unimodal) methods
2. Framework for learning from both demos + language in training + testing
3. Benchmark of 300 highly-randomized pick-and-place tasks

Main Takeaways

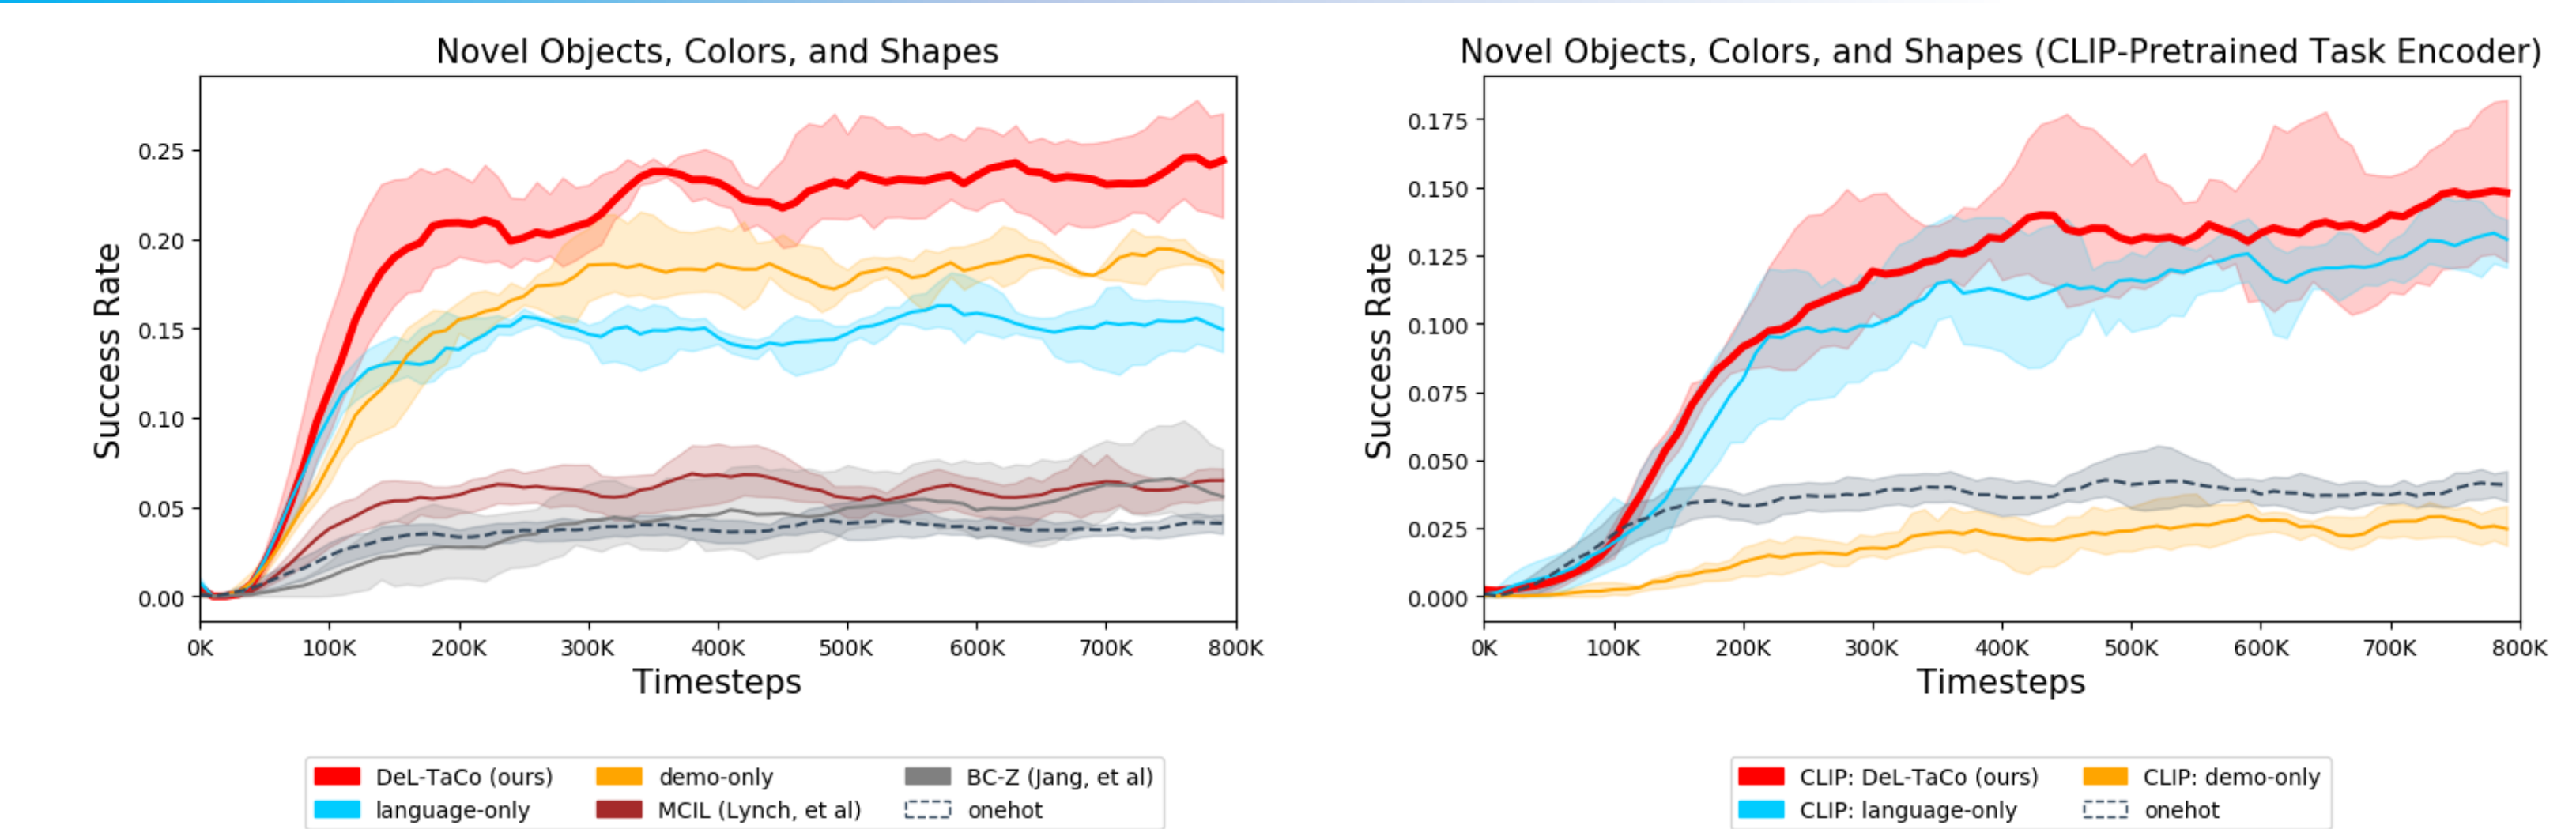
- 1. Generalization:** Learning from both demonstrations and language with DeL-TaCo leads to better generalization than learning from one modality
- 2. Sample-Efficiency:** DeL-TaCo reduces the need for ~50 demonstrations over policies conditioned only on demonstrations

Architecture



Results

Generalization to Novel Objects, Colors, and Shapes



Demo Encoder	Language Encoder	Task Conditioning	Success Rate ± SD (%)
-	-	One-hot (lower bound)	4.9 ± 1.7
-	-	One-hot Oracle (upper bound)	69.3 ± 7.4
CLIP (frozen)			
-	-	Language-only	14.4 ± 2.4
-	-	Demo-only	3.5 ± 1.0
-	-	DeL-TaCo (ours)	15.8 ± 2.2
-	miniLM (frozen)	Language-only	17.1 ± 2.2
CNN	-	Demo-only	20.8 ± 2.4
CNN	-	BC-Z [1] (Demo-only)	6.7 ± 2.3
CNN	miniLM (frozen)	MCIL [2] (Demo-only or Lang-only)	7.5 ± 1.2
CNN	miniLM (frozen)	DeL-TaCo (ours)	25.8 ± 3.4

[1] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, C. Finn. "BC-Z: Zero-Shot Task Generalization with Robotic Imitation Learning." CoRL, 2021.
 [2] Corey Lynch and Pierre Sermanet. "Language Conditioned Imitation Learning Over Unstructured Data." RSS, 2021.

How many Demonstrations is Language Learning Worth?

Task Conditioning	Demo-only					DeL-TaCo (ours)
# demos per test-task finetuned on	0	10	25	50	100	0
Success Rate (%)	20.8	23.4	24.6	26.1	32.9	25.8
± SD (%)	±2.4	±1.8	±2.5	±2.6	±2.5	±3.4