



LISA: Learning Interpretable Skill Abstractions from Language

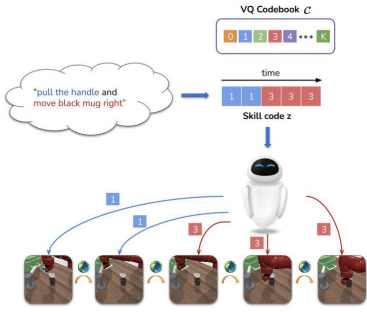
Divyansh Garg



SCAN ME

LISA

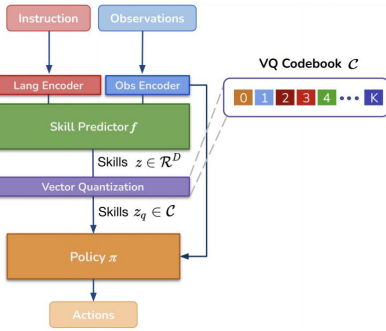
Goal: Enable agents to follow language instructions in complex multi-task environments




LISA: Learn interpretable high-level skills + low-level policy

- ✓ Skill codes are quantized and interpretable
- ✓ Skill codes can be composed to get desired behavior
- ✓ Generalize to very long-range & unseen tasks
- ✓ Works very well in the low-data regime

Architecture



Interpretable Skill Codes



Skill Heat Maps






Figure 3: Behavior with fixed LISA options. We show the word clouds and the behavior of the policy obtained by using a fixed skill code $z = 14$ for an entire episode. We find that this code encodes the skill "closing the drawer", as indicated by the word cloud. The policy executes this skill with a high degree of success when conditioned on this code for the entire trajectory, across different environment initializations and seeds.

Environments




BabyAI



LORL

Learnt Skills

Figure 5: Word clouds on LORL. We show the most correlated words for 4 different learnt skill codes on LORL. We can see that the codes represent interpretable and distinguishable skills. For e.g., the code on the top left corresponds to closing the drawer. (note that container is a synonym for drawer in the LORL dataset)



Results

Table 1: Imitation Results: We show our success rates (in %) compared to the original method and a flat non-hierarchical baseline on each dataset. LISA outperforms all other methods in the low-data regime, and reaches similar performance as the number of demonstrations increases. Best method shown in bold.

Task	Num Demos	Original	Flat Baseline	LISA
BabyAI GoToSeq	1k	33.3 ± 1.3	49.3 ± 0.7	59.4 ± 0.9
BabyAI GoToSeq	10k	40.4 ± 1.2	62.1 ± 1.2	65.4 ± 1.6
BabyAI GoToSeq	100k	47.1 ± 1.1	74.1 ± 2.3	77.2 ± 1.7
BabyAI SynthSeq	1k	12.9 ± 1.2	42.3 ± 1.3	46.3 ± 1.2
BabyAI SynthSeq	10k	32.0 ± 2.5	52.1 ± 0.5	53.3 ± 0.7
BabyAI SynthSeq	100k	40.4 ± 3.3	64.2 ± 1.3	61.2 ± 0.6
BabyAI BossLevel	1k	20.7 ± 4.6	41.5 ± 3.3	49.1 ± 2.4
BabyAI BossLevel	10k	28.9 ± 1.3	60.1 ± 5.5	58 ± 4.1
BabyAI BossLevel	100k	45.3 ± 0.9	72.0 ± 4.2	69.8 ± 3.1
LORL - States (fully obs.)	50k	6 ± 1.2 [†]	33.3 ± 5.6	66.7 ± 5.2
LORL - Images (partial obs.)	50k	29.5 ± 0.07 [†]	15 ± 3.4	40 ± 2.0

Outperform non-hierarchical methods in low-data regime

Composition Tasks

Table 2: LISA Composition Results: We show our performance on the LORL Sawyer environment compared to baselines

Method	Success Rate (in %)
Flat	8.33
LORL Planner	20.1
LISA (Ours)	22.5

Test on long composition instructions like "close the drawer, turn the faucet left and move black mug right"

Outperform prior methods by more than 2x