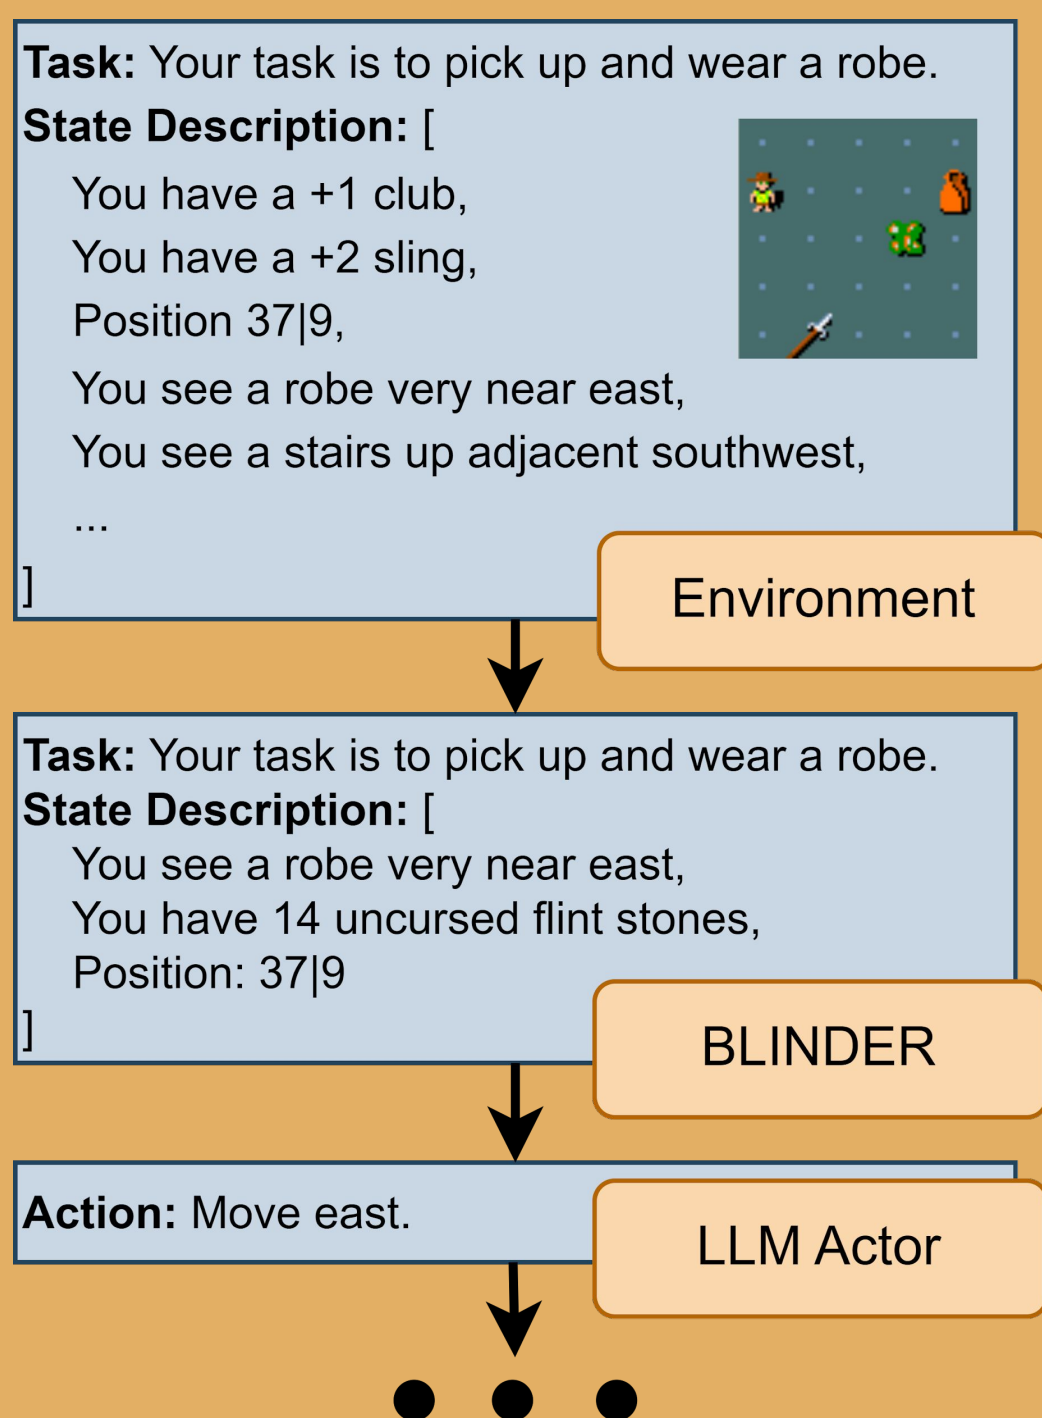# Selective Perception: Learning Concise State Descriptions for Language Model Actors

Kolby Nottingham - Yasaman Razeghi - Kyungmin Kim - JB Lanier
Pierre Baldi - Roy Fox - Sameer Singh
University of California Irvine

Motivation

Task: Your task is to pick up and wear a robe.
State Description: [
  You have a +1 club,
  You have a +2 sling,
  Position 37|9,
  You see a robe very near east,
  You see a stairs up adjacent southwest,
  ...
]

Environment

Task: Your task is to pick up and wear a robe.
State Description: [
  You see a robe very near east,
  You have 14 uncursed flint stones,
  Position: 37|9
]

BLINDER

Action: Move east.

LLM Actor

- LLMs are used for many embodied tasks with complex state spaces
- State descriptions often include unnecessary or distracting information
- Concise descriptions can lower cost and improve performance

Can we learn to select only the most relevant parts of a state description?

Method

BLINDER Training
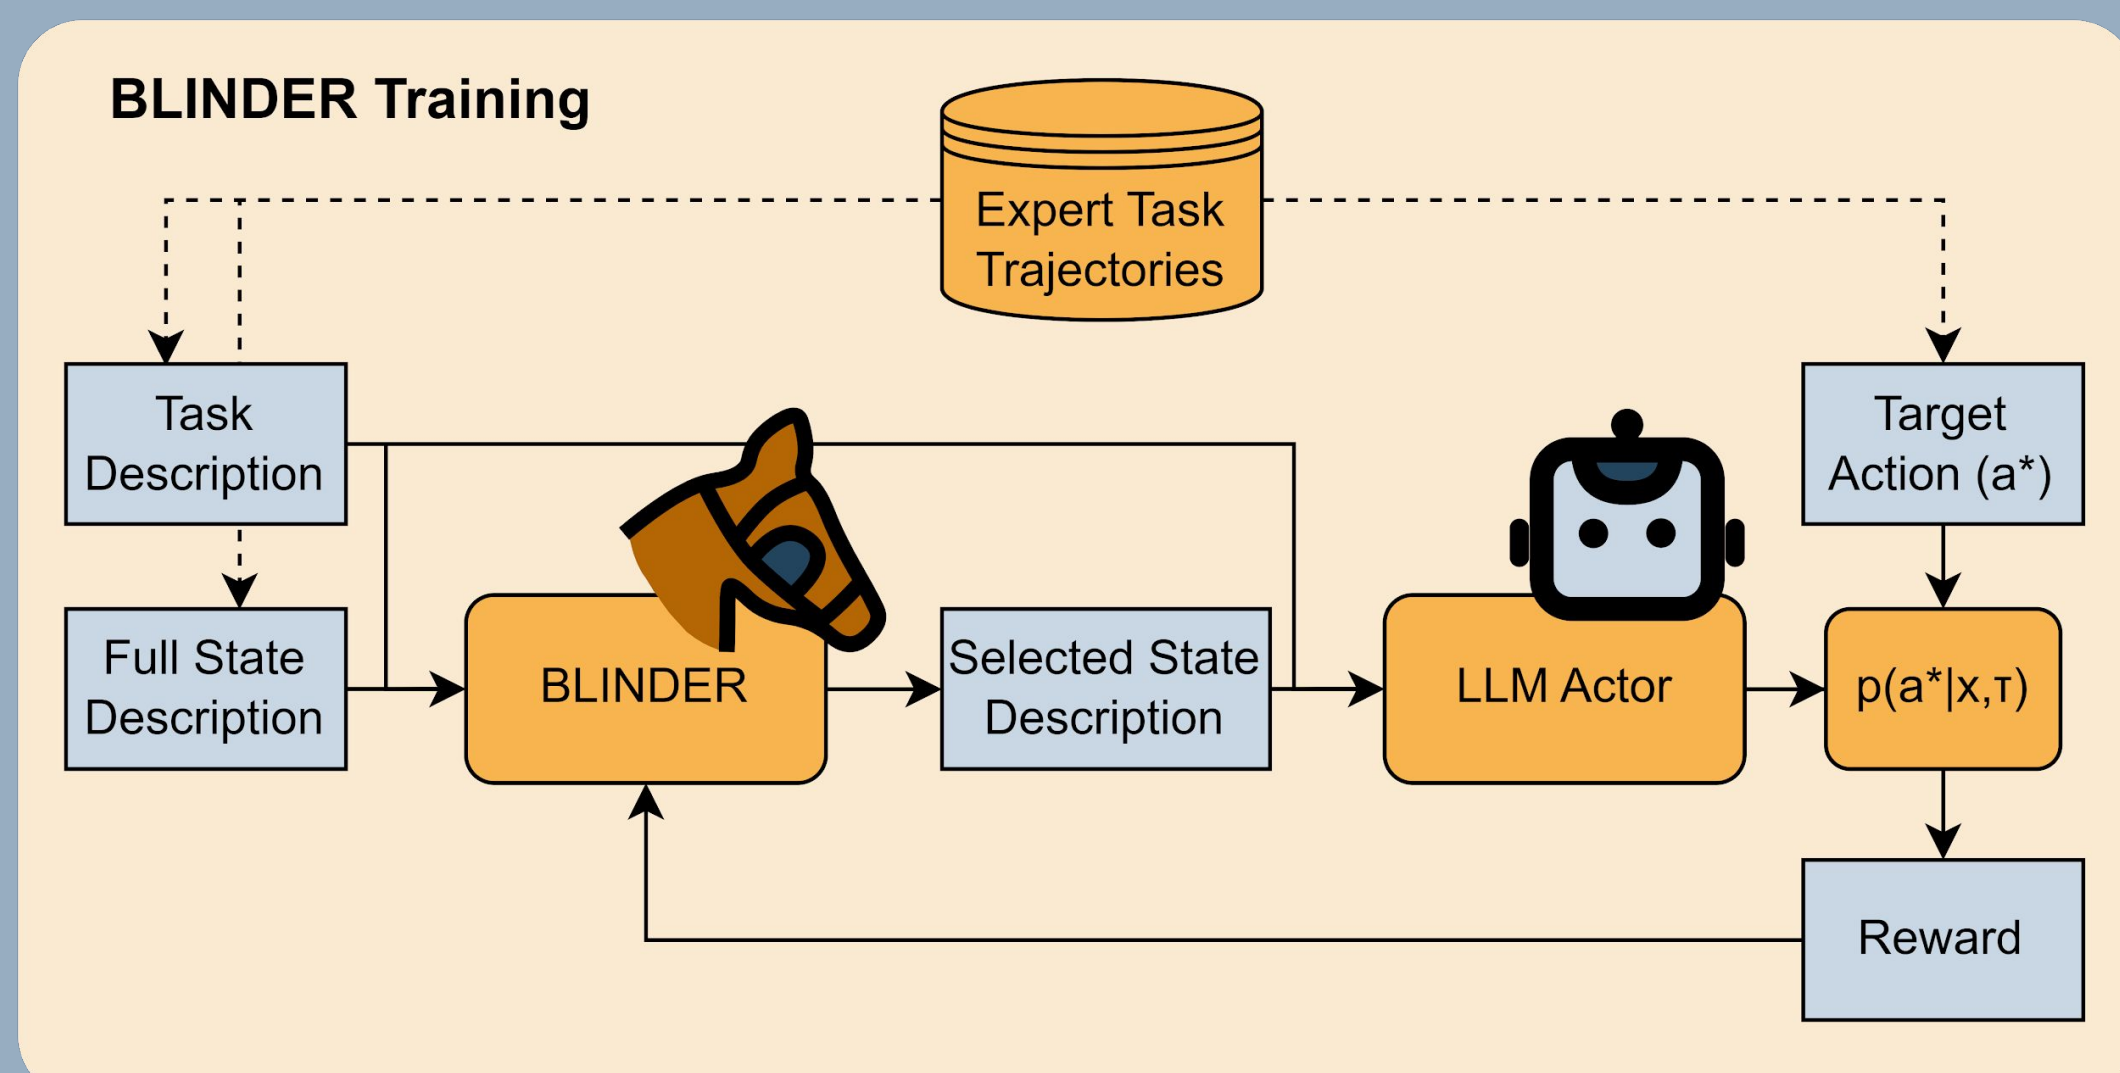
1. Demonstrate a handful of training tasks (~5)
2. Sample state descriptions from demonstrations
3. Use likelihood of demonstrated actions from an LLM actor as a reward signal
4. Finetune an LLM value function via value-based RL
5. Use the value function to sample state descriptions at inference time
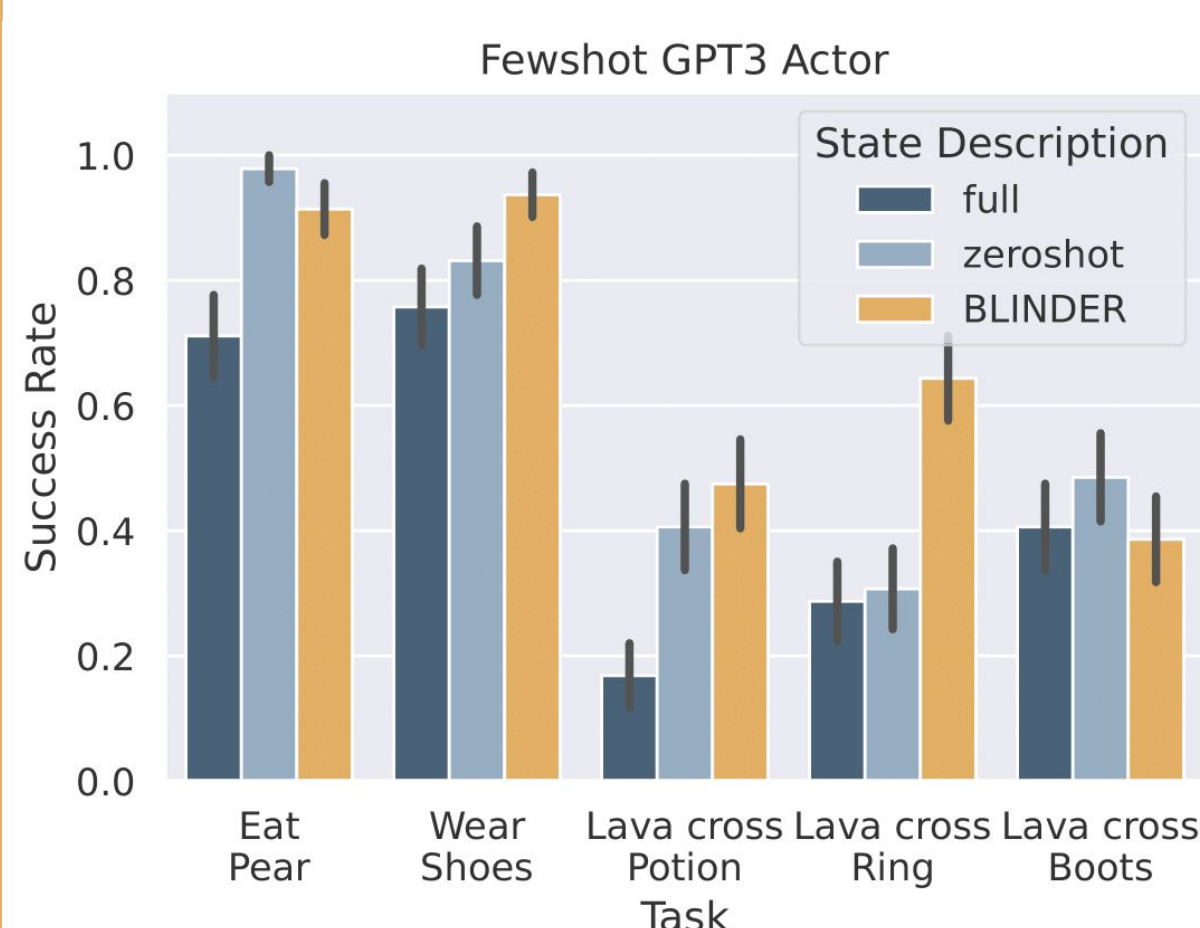
NetHack Experiments
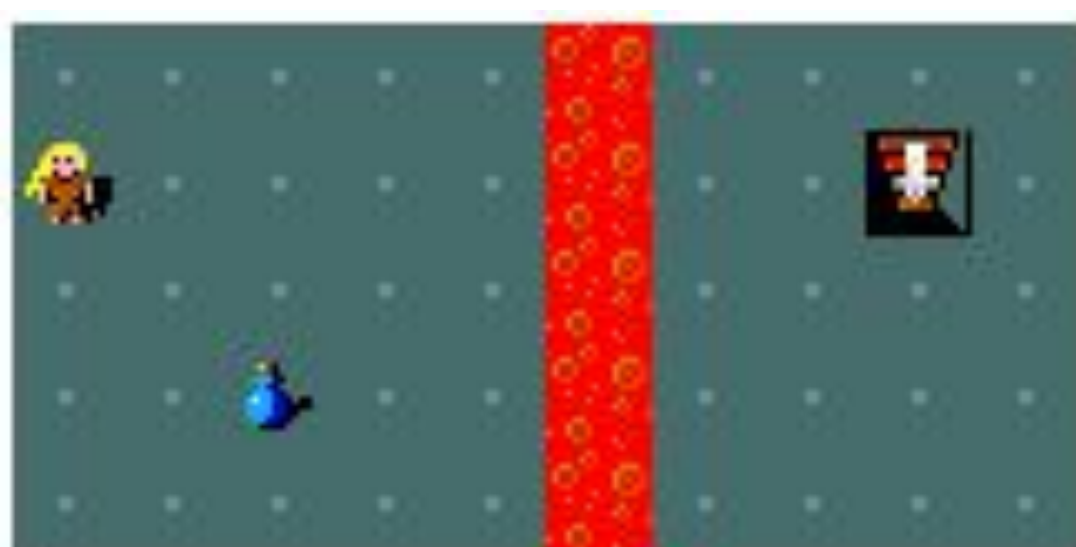
NetHack is a game in which the agent navigates to the next staircase. The state is described as a list of objects with their direction and distance.
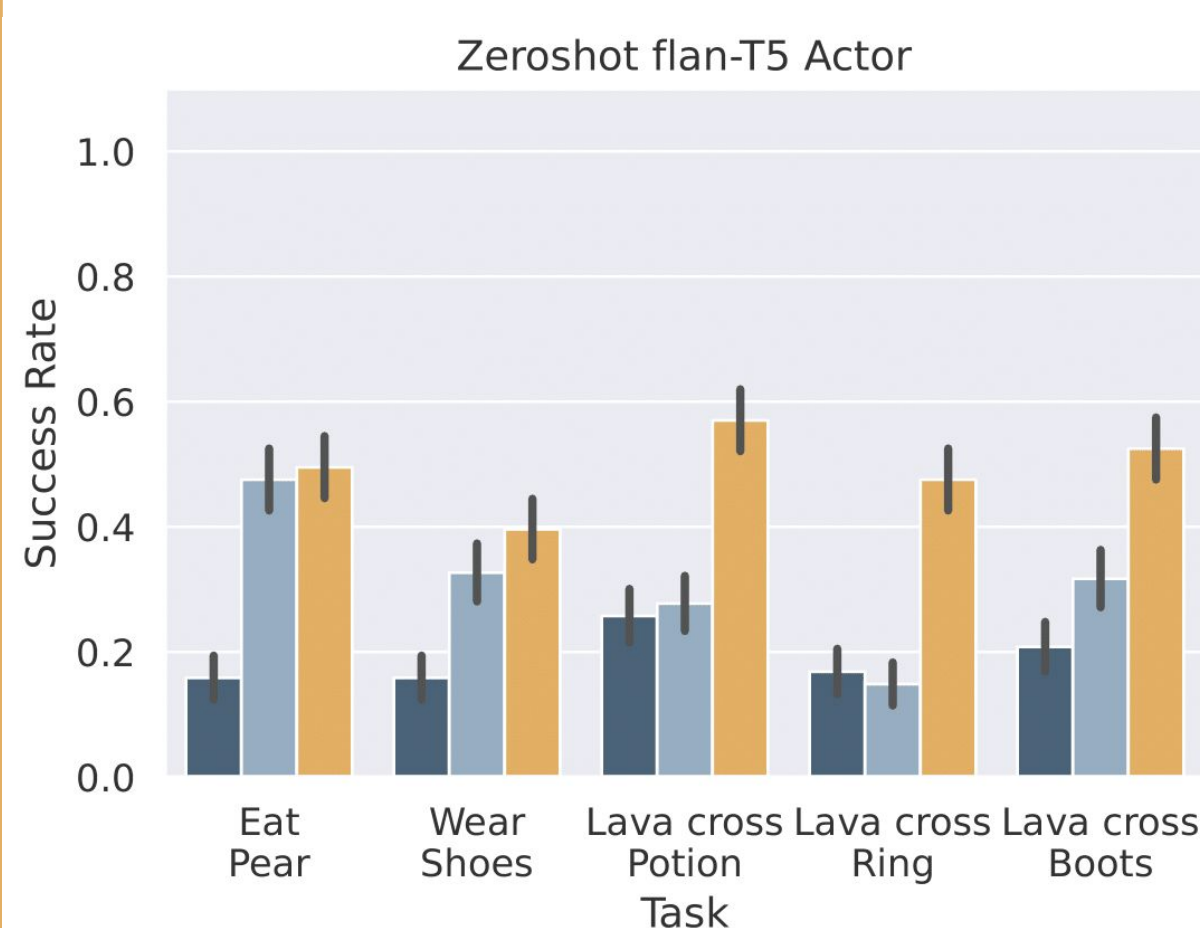
BLINDER outperforms exhaustive descriptions and zeroshot summaries from similar sized models
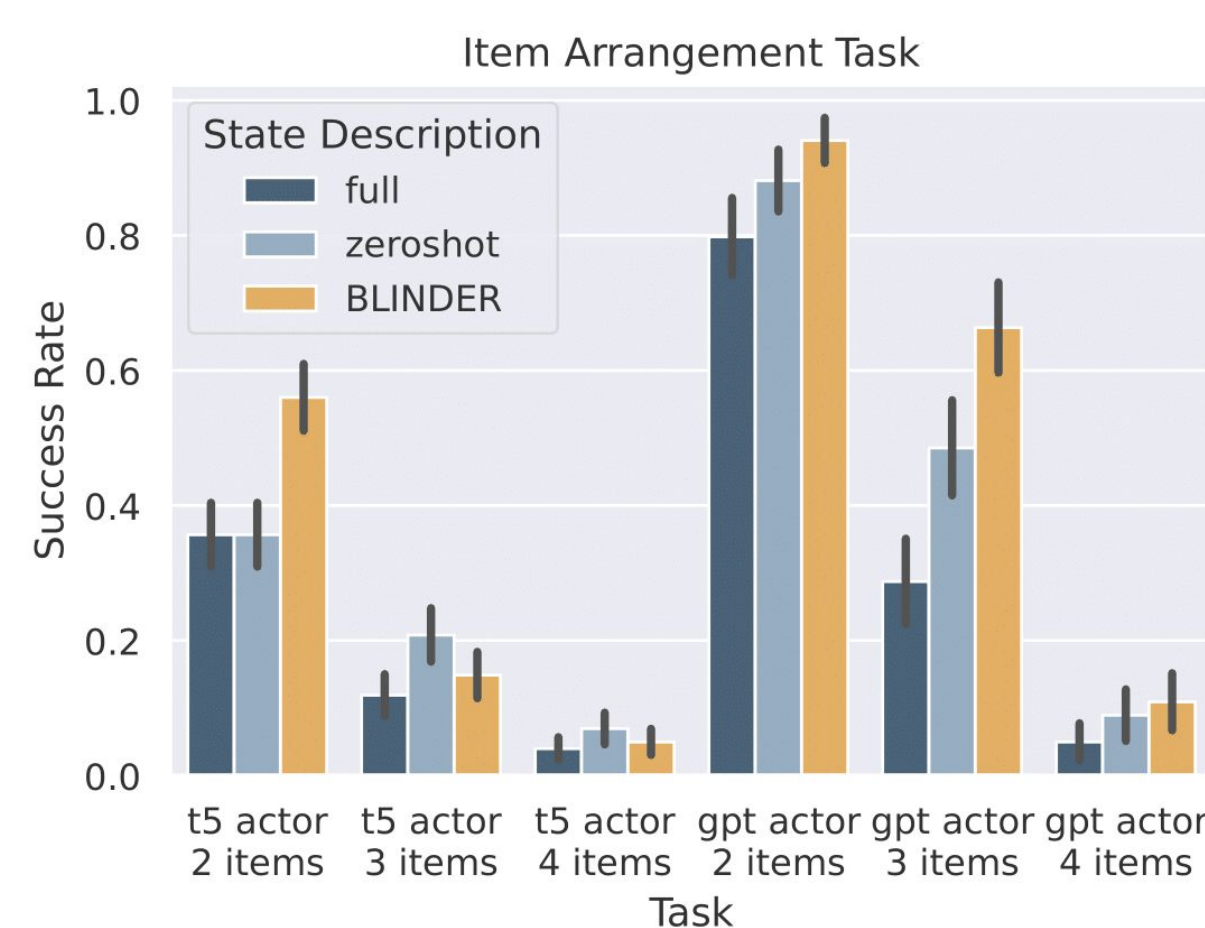
BLINDER generalizes to larger actors, regularly beating baselines at 1/6th the context length

Robotic Experiments

The LLM actor directs a robot to sort objects on a table from left to right. The state is described as the relative positions of all objects to each other.

BLINDER continues to beat baselines by selecting the most relevant parts of the state description for the LLM actor.

https://kolbytn.github.io/blinder/

knotting@uci.edu

UCI