Scaling Up and Distilling Down
Language-Guided Robot Skill Acquisition

Huy Ha
Columbia University

Pete Florence
Google DeepMind

Shuran Song
Columbia University

Language-guided Robot Skill Learning

For SCALING UP data generation, a language model guides sampling-based robot planners to generate **rich and diverse** manipulation trajectories.

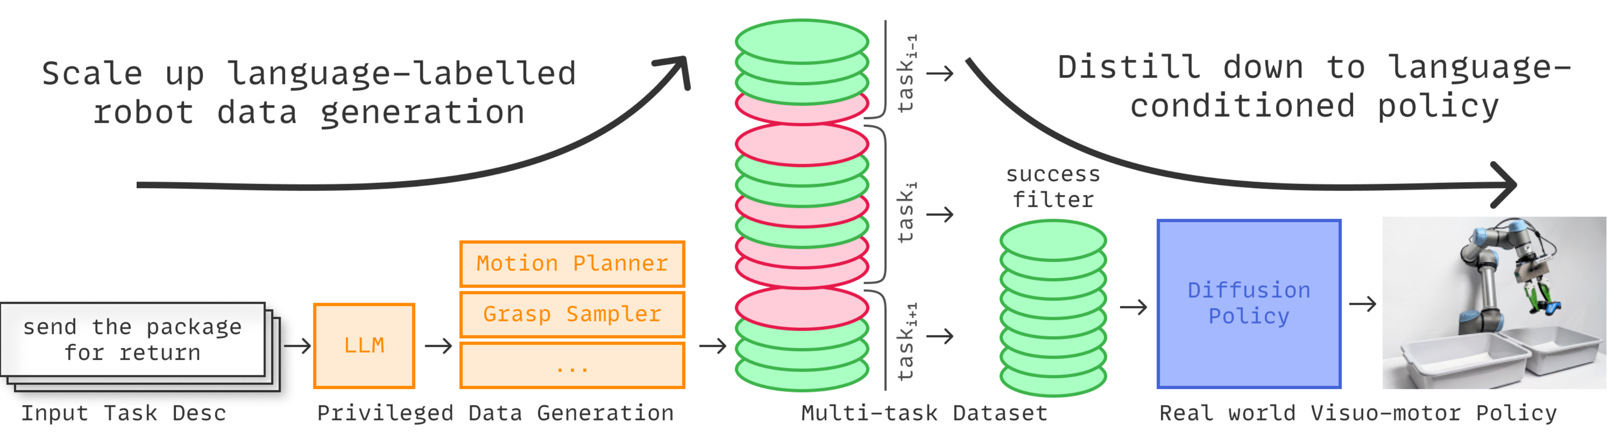

For DISTILLING DOWN for real world deployment, we extend Diffusion Policy to multi-task settings with language conditioning.

Language-guided not Language-constrained

Our framework is a step towards putting robotics on the same scaling trend as large language models while *not compromising* on rich low-level manipulation.

Approach	6 DoF Manipulation	Common-sense	No Sim State
Sampling-based Planners	✓	✗	✗
LLM Planners	✗	✓	✗
Our Data Generation	✓	✓	✗
Our Policy	✓	✓	✓

Run it in Real with Zero Finetuing

Using domain randomization, we deployed our diffusion policy on a real robot with no fine-tuning.

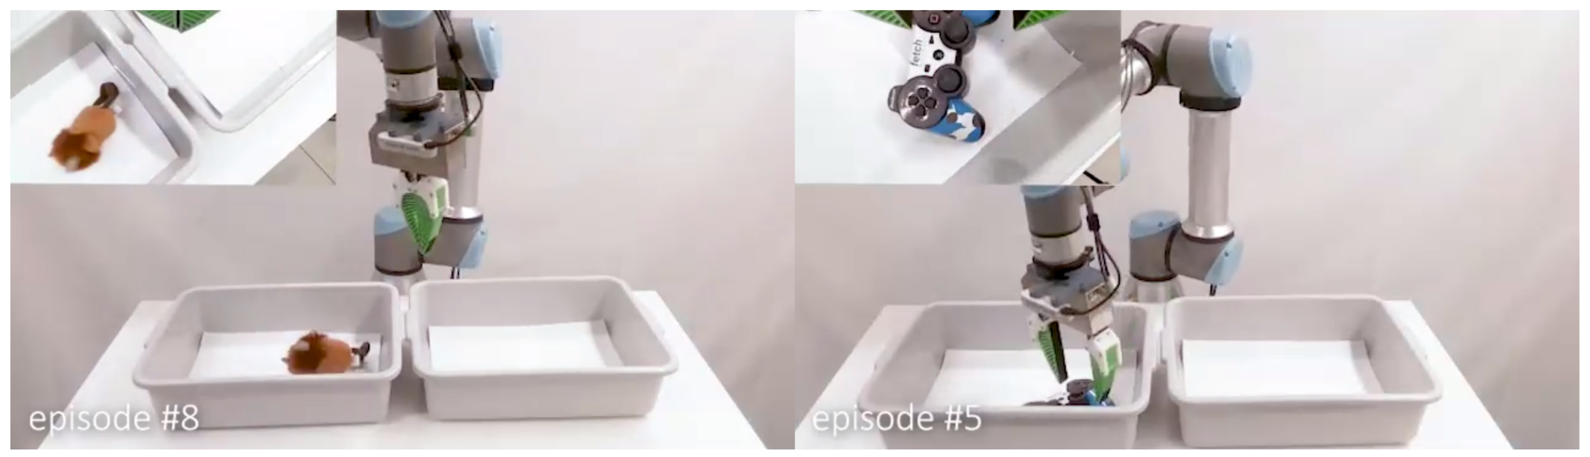

Scaling Up and Distilling Down is a framework for language-guided skill learning. Give it a task description, and it will automatically generate rich, diverse robot trajectories, complete with success label and dense language labels.

The best part? It uses no expert demonstrations, manual reward supervision, and no manual language annotation.

A New Multi-Task Benchmark

DOMAINS	TASKS	ASSETS	LANGUAGE LABELS
5	18	30	∞

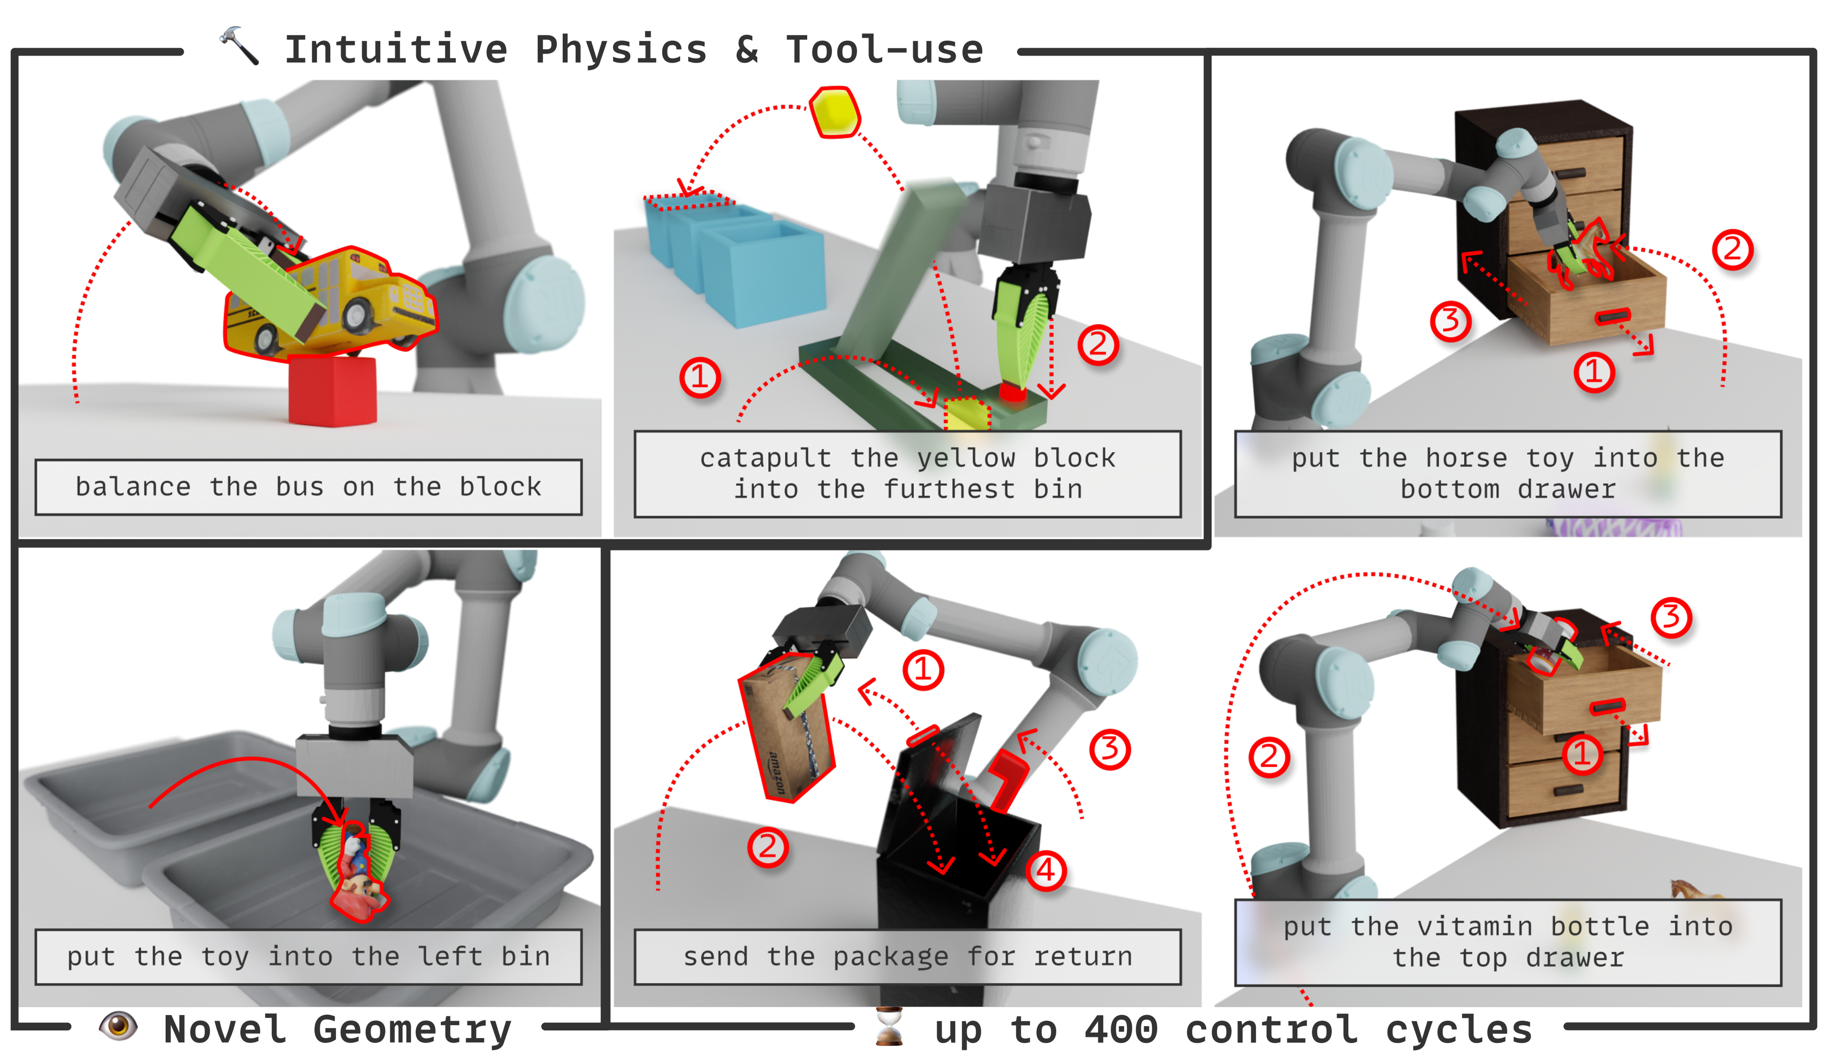

Language-conditioned Diffusion Policies

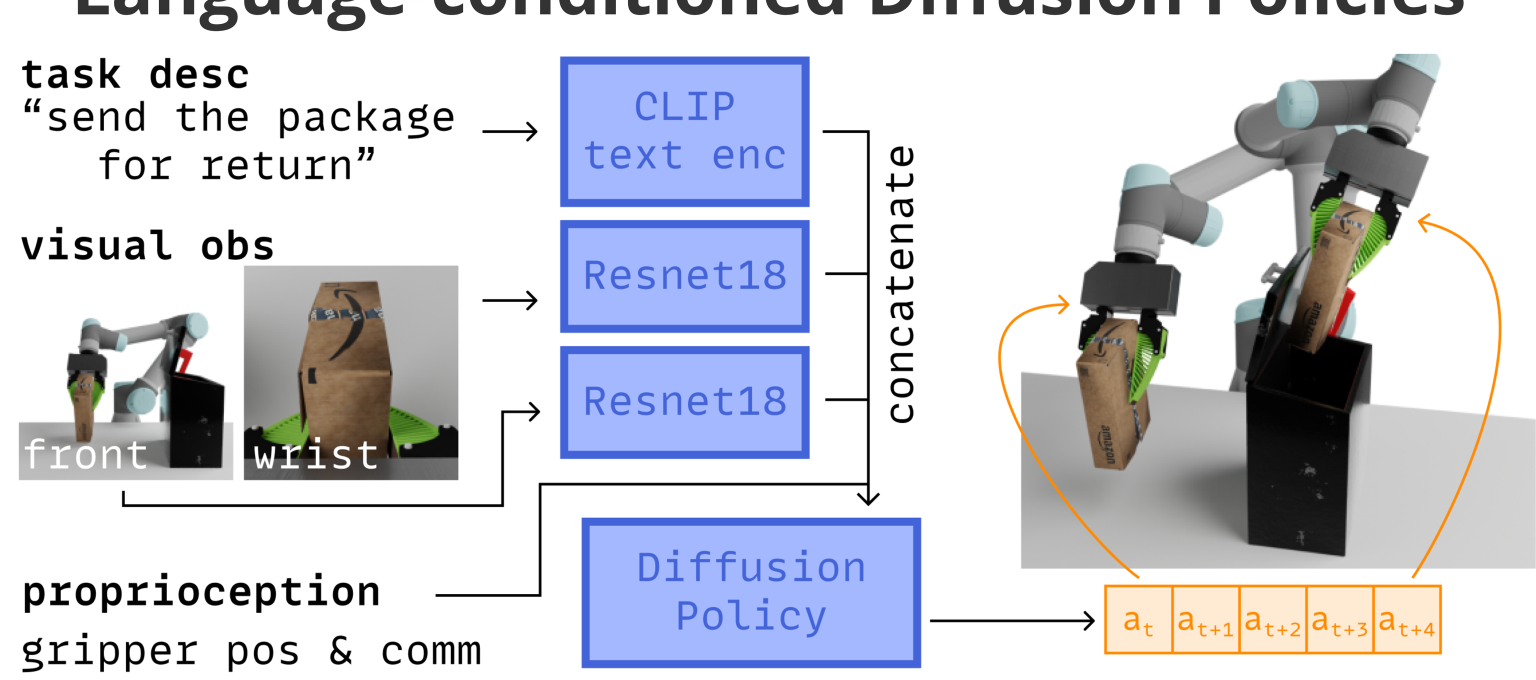

Robust Retries After Failures

Using a language model to predict each task's success condition code snippet allows the robot to retry failed tasks. The result are demonstrations of robust behavior, which teach the policy to **recover after failed attempts.**

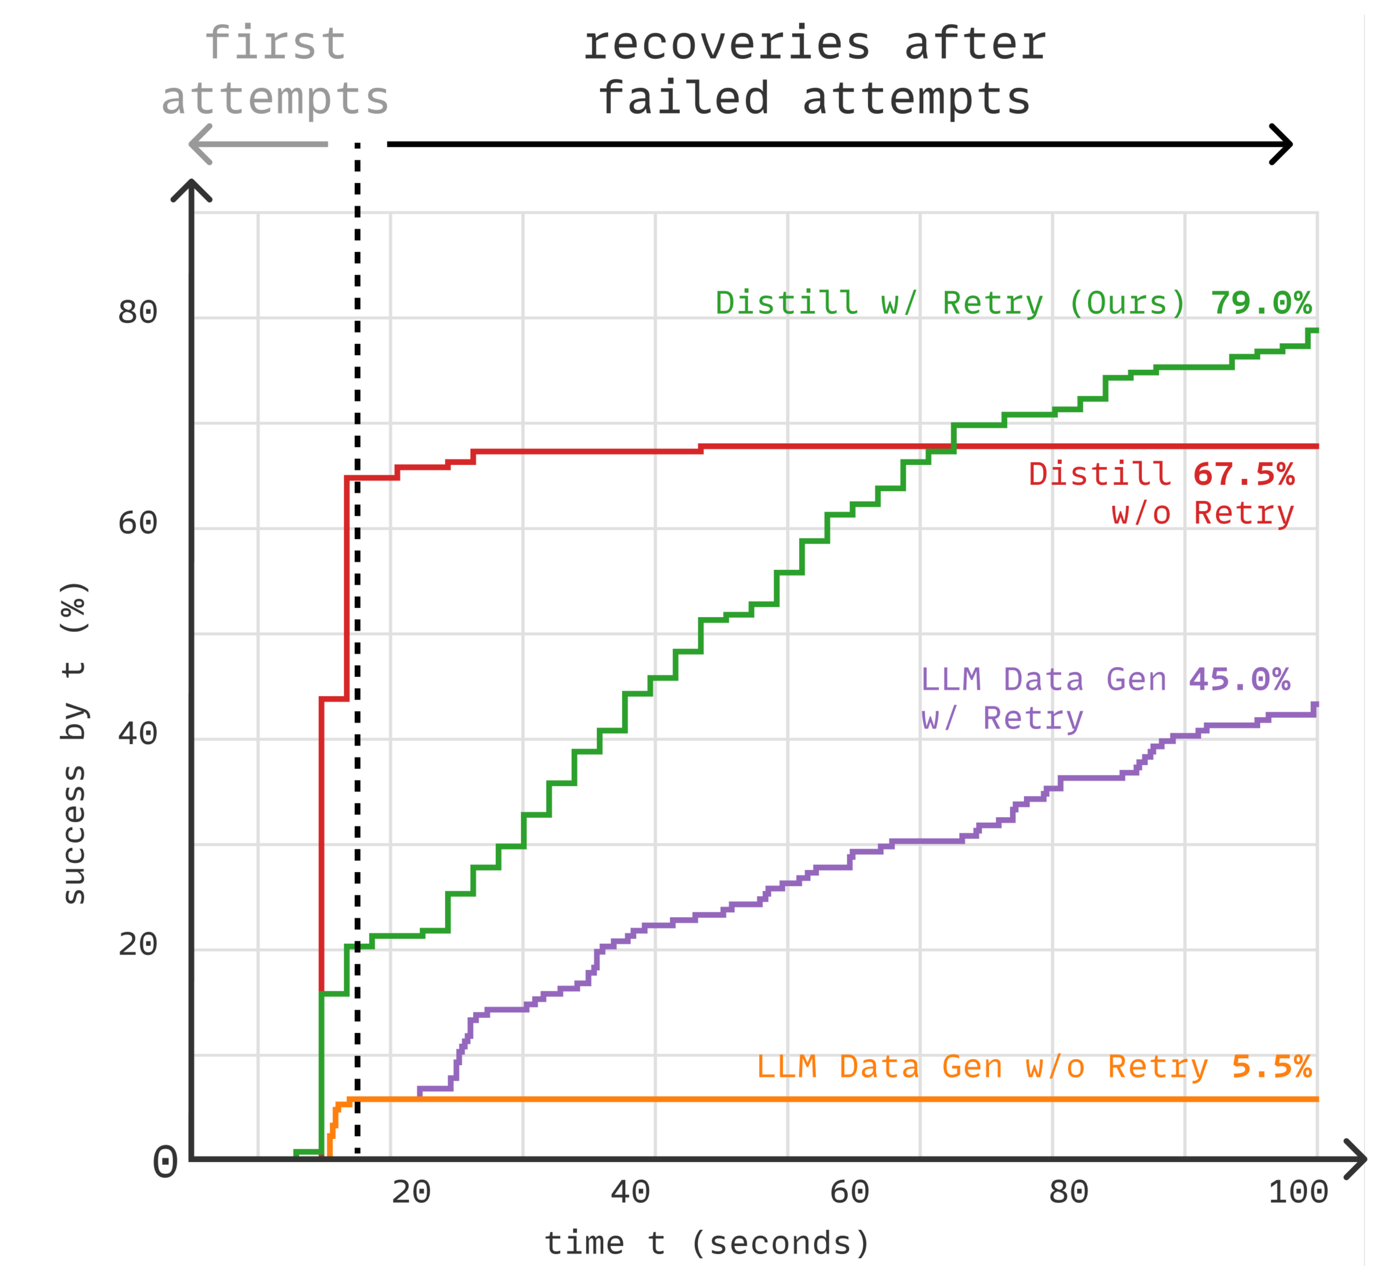

See how well your policy scales with infinite language-labeled and diverse robot trajectories.