Ambient Diffusion Policy

Imitation Learning From Suboptimal Data

Adam Wei, Nicholas Pfaff, Thomas Cohn, Arif Kerem Dayı, Constantinos Daskalakis, Giannis Daras, Russ Tedrake

Method

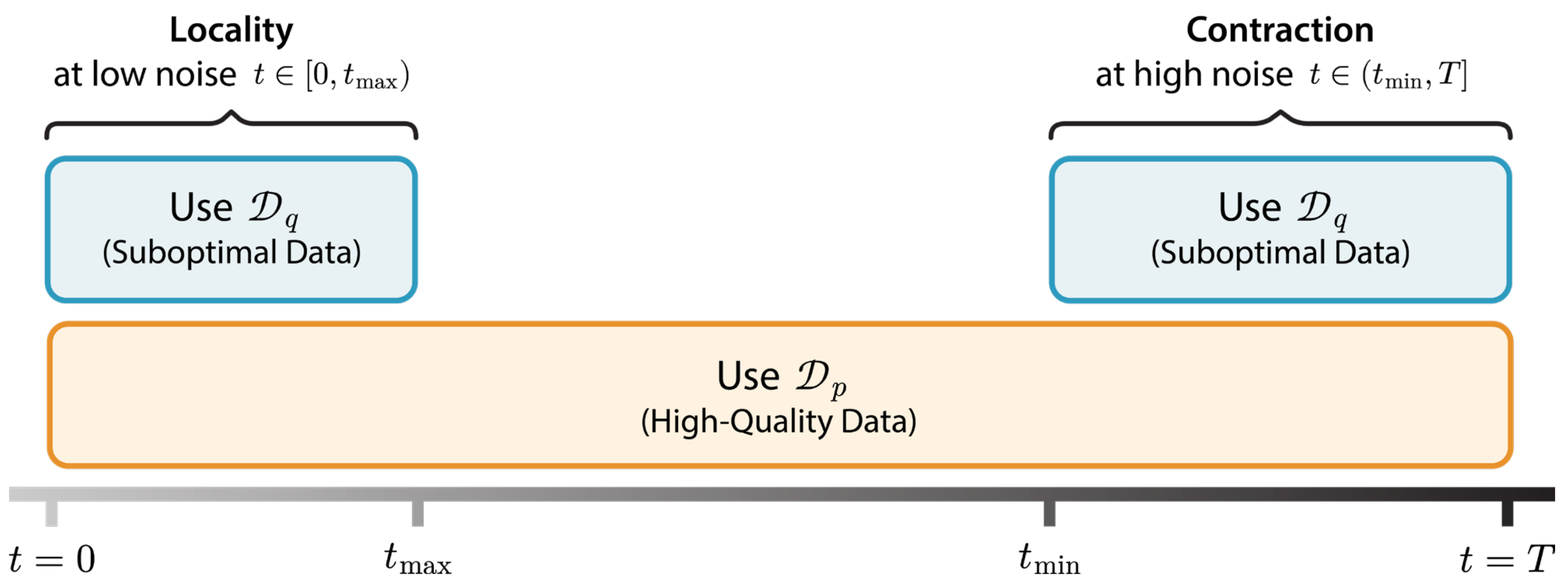

"Suboptimal" data is everywhere in robotics

- Low-quality teleop
- Ego-centric data
- Simulation
- Incorrect task
- Cross-embodied
- ... and many more

Ambient Diffusion Policy restricts "suboptimal" data to training at ***high*** and ***low*** noise-levels

Global-to-Local Hierachy

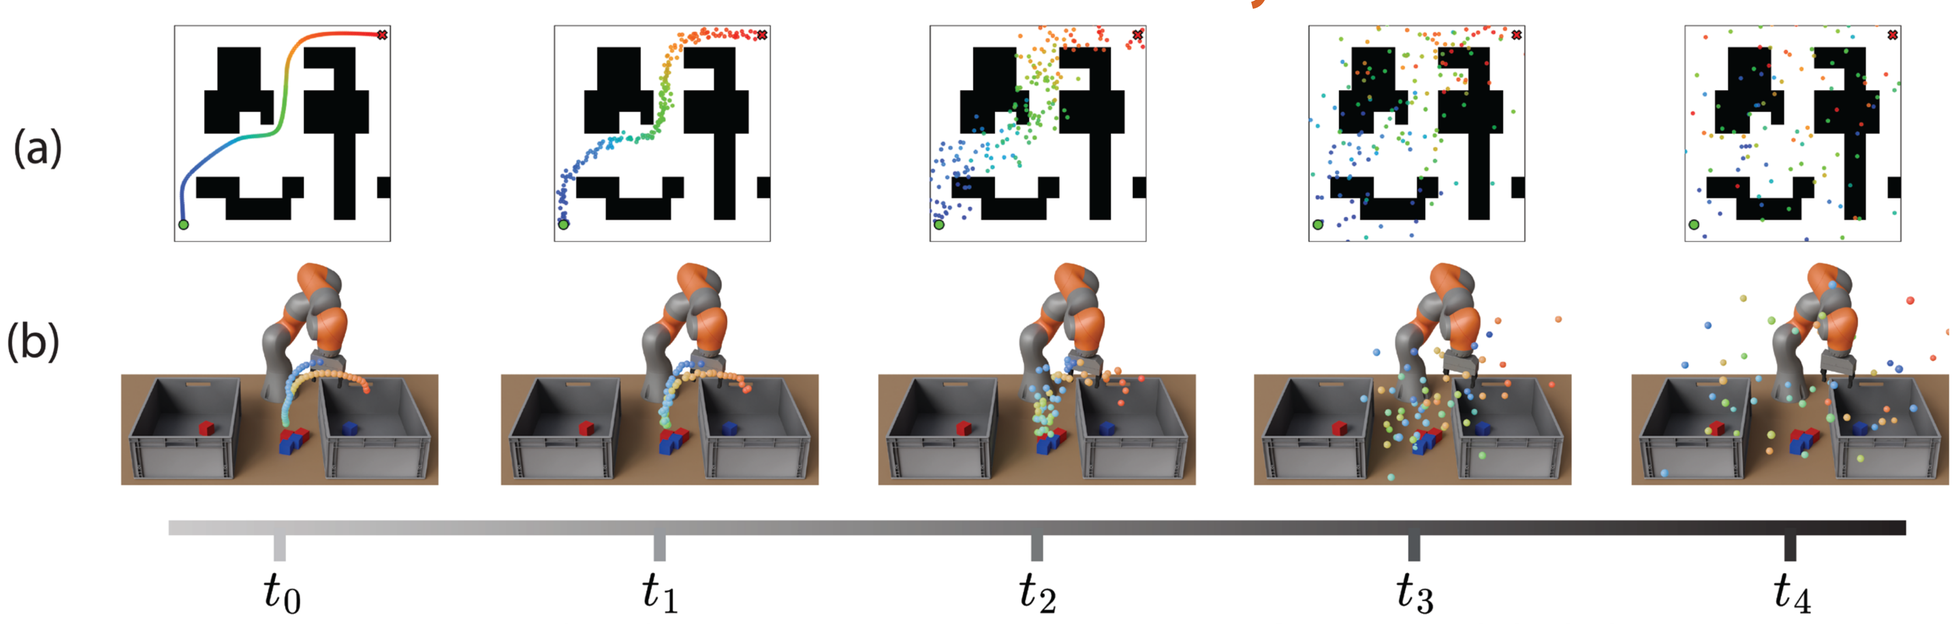

Low noise: *local* motion primitives **High noise:** *global* planning

Diffusion Policy learns different features at each noise-level $\Longrightarrow$ **data usage should be noise-dependent**

Ambient + Open X-Embodiment

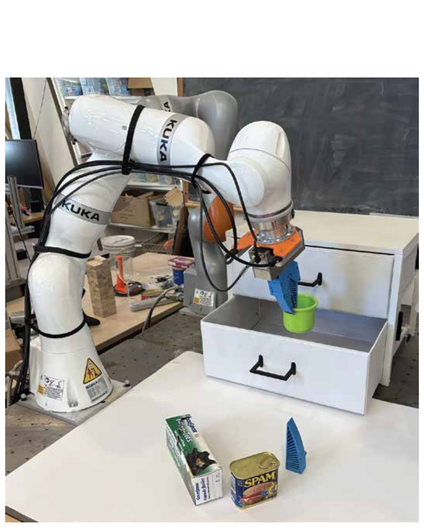

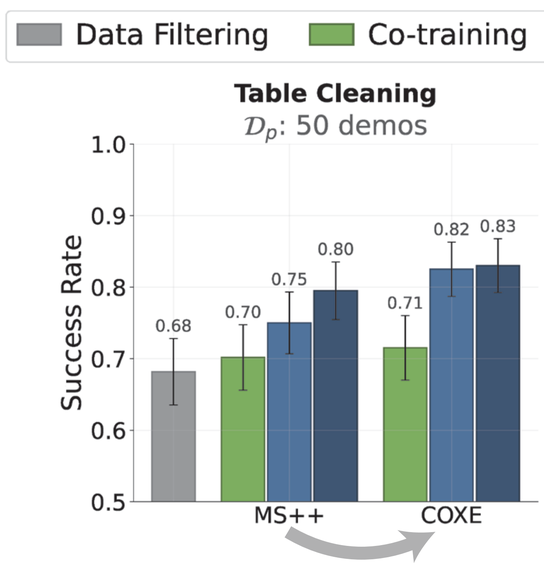

1) Ambient ***scales with more suboptimal data***, whereas co-training plateaus

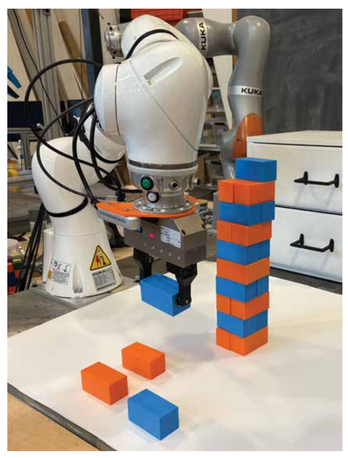

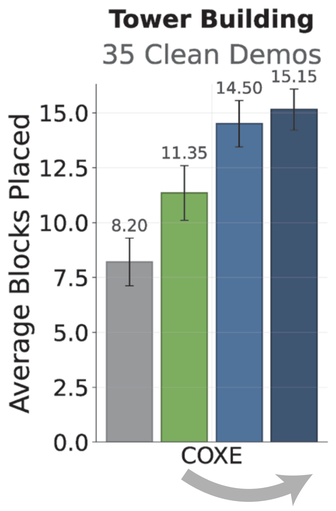

2) Ambient ***outperforms*** co-training by up to 33%

Generality

Noisy Teleop
Neural Motion Planning

Sim-and-Real
Planar Pushing

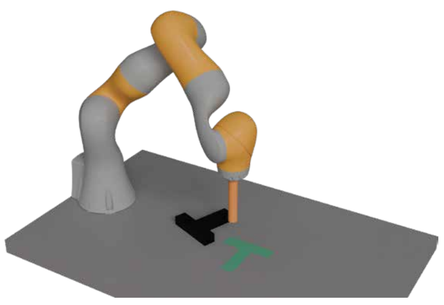

Task Mistatch
Block Sorting

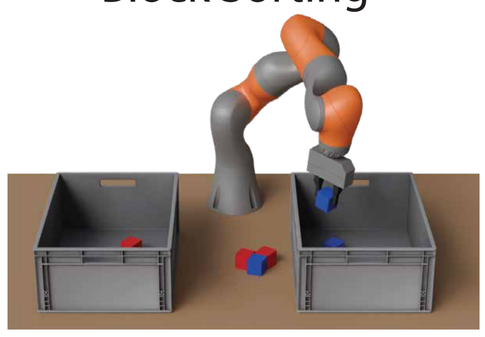

	Noisy Teleop	Sim-and-Real	Task Mistatch
Co-training:	59.9% & Jittery	84.5%	22.7%
Ambient:	**65.9% & Smooth**	**93.5%**	**93.3%**

Learn More

Project Website