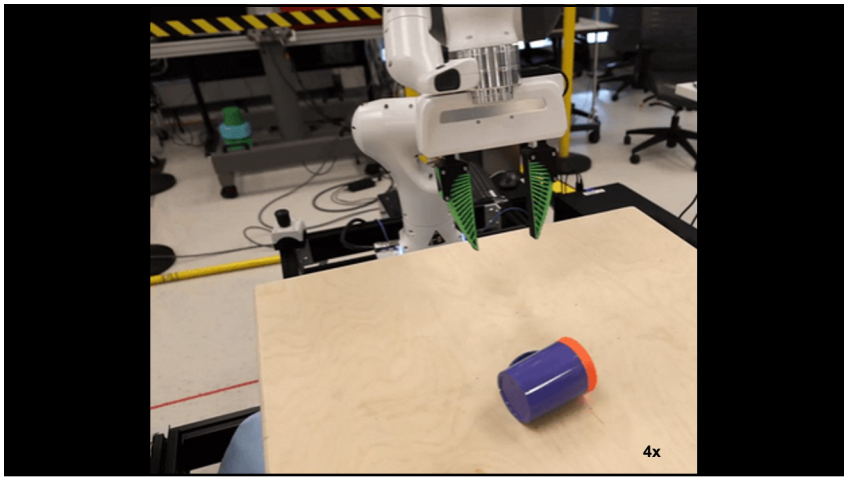
4x

1

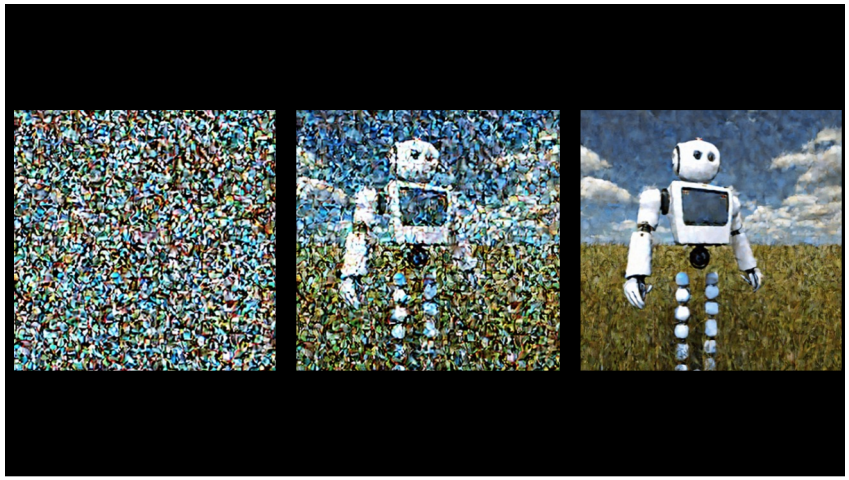

2

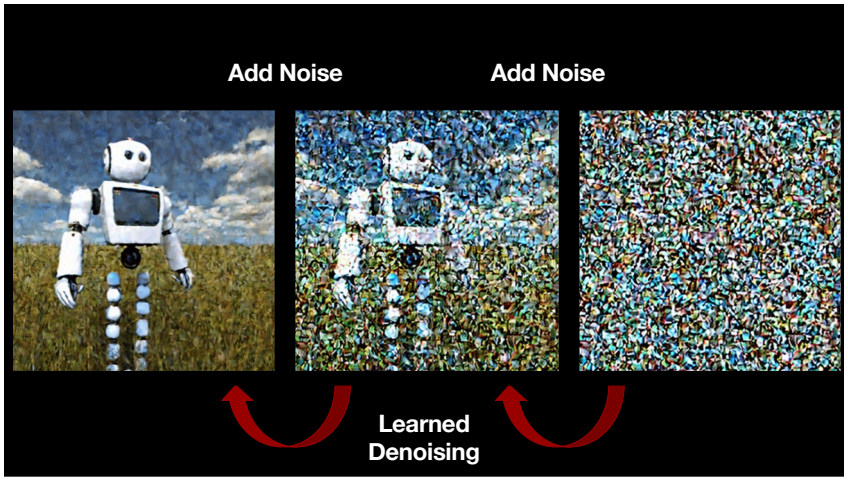
Add Noise
Add Noise
Learned
Denoising

3

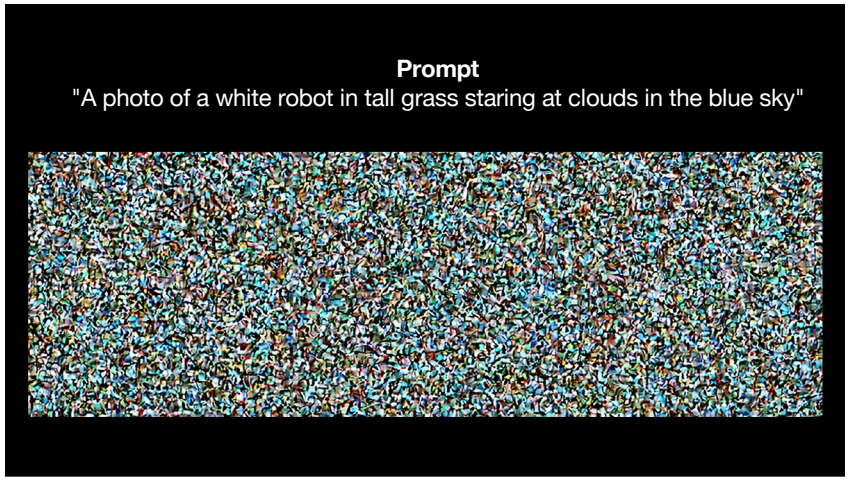
Prompt
"A photo of a white robot in tall grass staring at clouds in the blue sky"

4

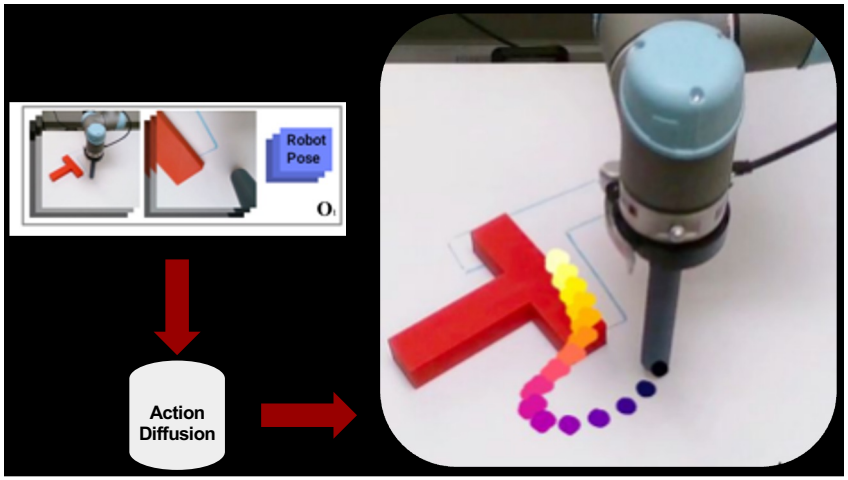

5

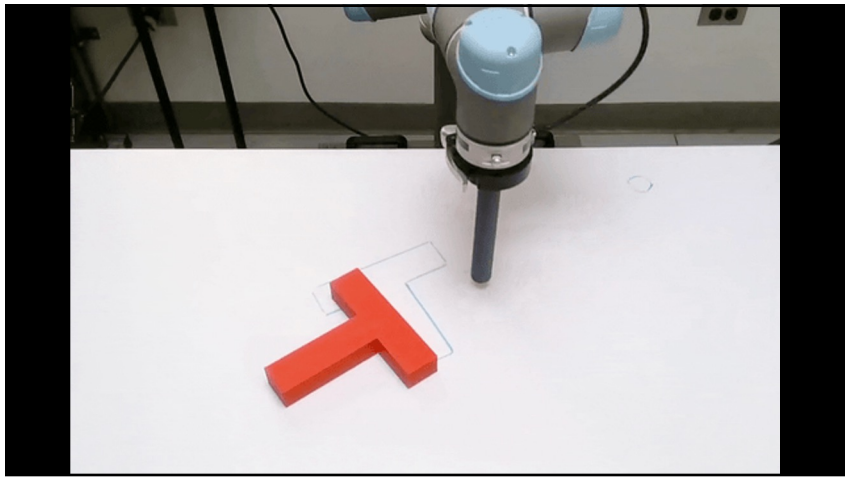

6

Diffusion Policies vs Other Representations

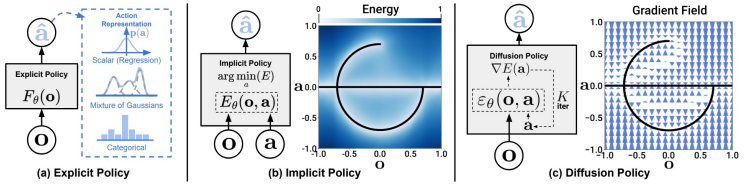

Fig. 1: **Policy Representations.** a) Explicit policy with different types of action representations. b) Implicit policy learns an energy function conditioned on both action and observation and optimizes for actions that minimize the energy landscape c) Diffusion policy refines noise into actions via a learned gradient field. This formulation provides stable training, allows the learned policy to accurately model multimodal action distributions, and accommodates high-dimensional action sequences.

7

Diffusion Policy
A General Recipe for Behavior Cloning

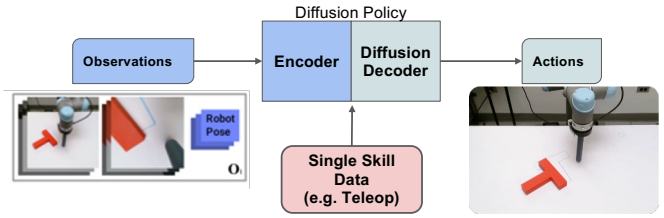

© 2023 Toyota Research Institute. Confidential. Do not distribute.

Diffusion Policy: Visuomotor Policy Learning via Action Diffusion
Chi et al. 2023

8

Original Single-Skill Architecture

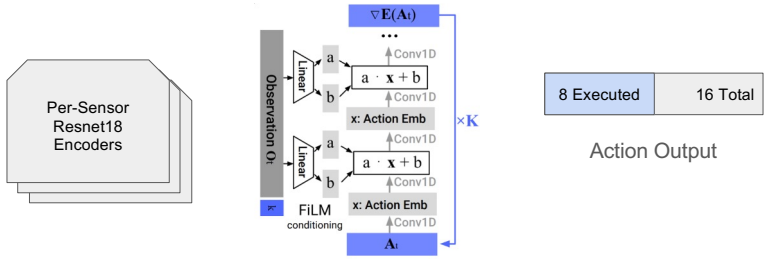

1D U-net Diffusion Head

9

Why Diffusion?

10

Performance

Training Stability and Consistency

Multimodality

11

Performance

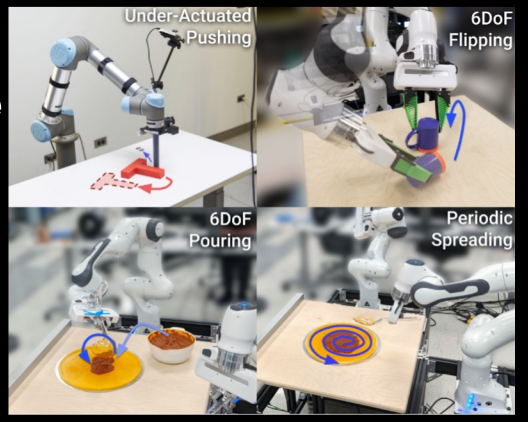

12

13

14

Consistency

Hyperparameter Tuning

15

Multimodality

16

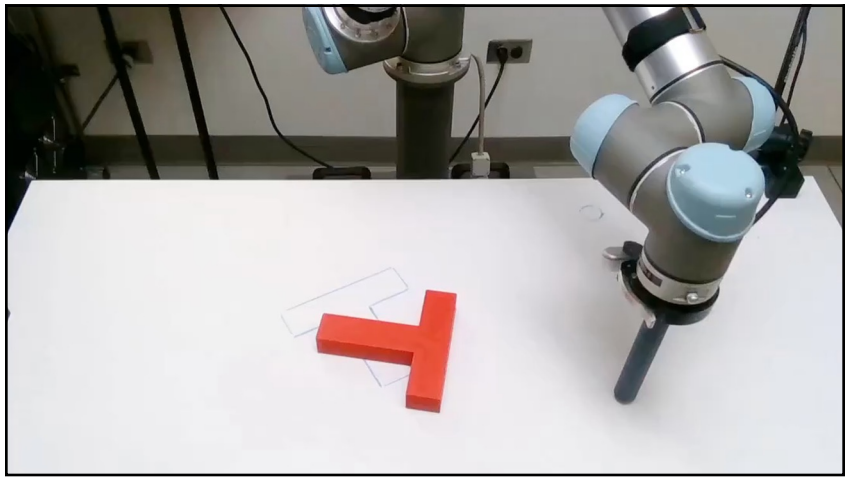

17

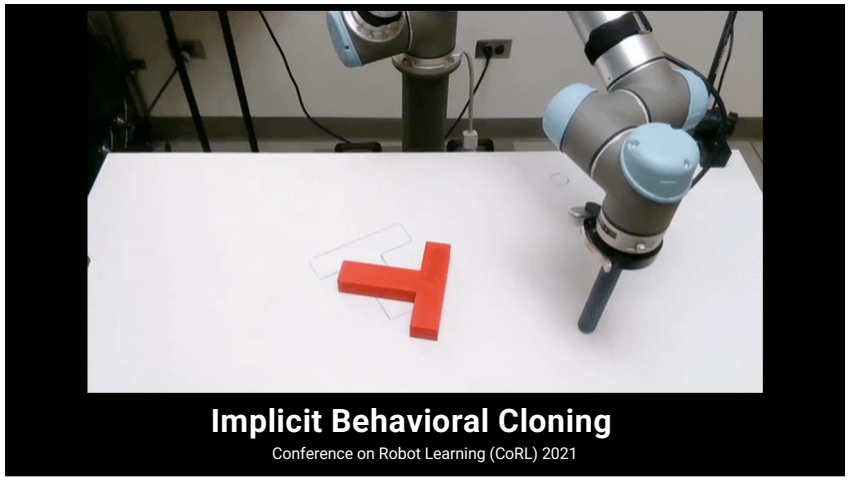
Implicit Behavioral Cloning
Conference on Robot Learning (CoRL) 2021

18

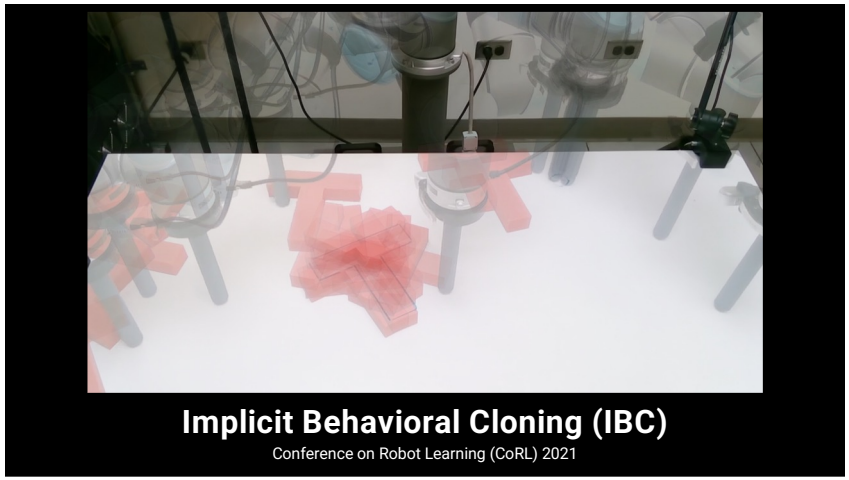
Implicit Behavioral Cloning (IBC)
Conference on Robot Learning (CoRL) 2021

19

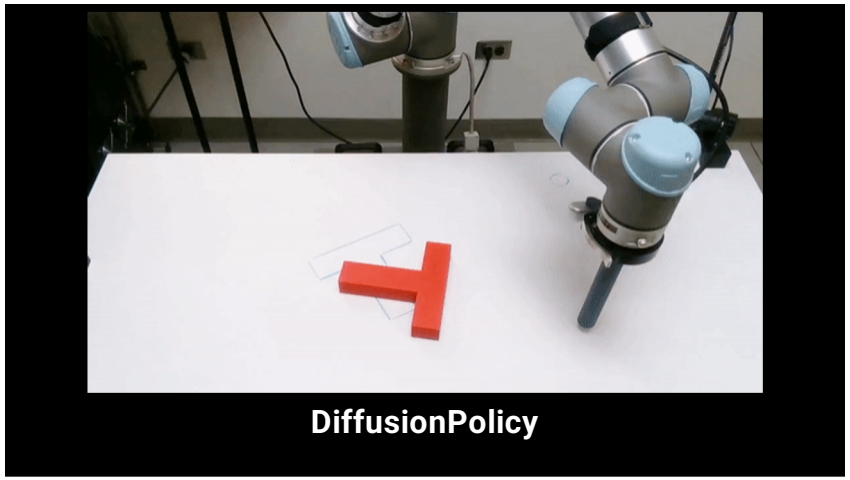
DiffusionPolicy

20

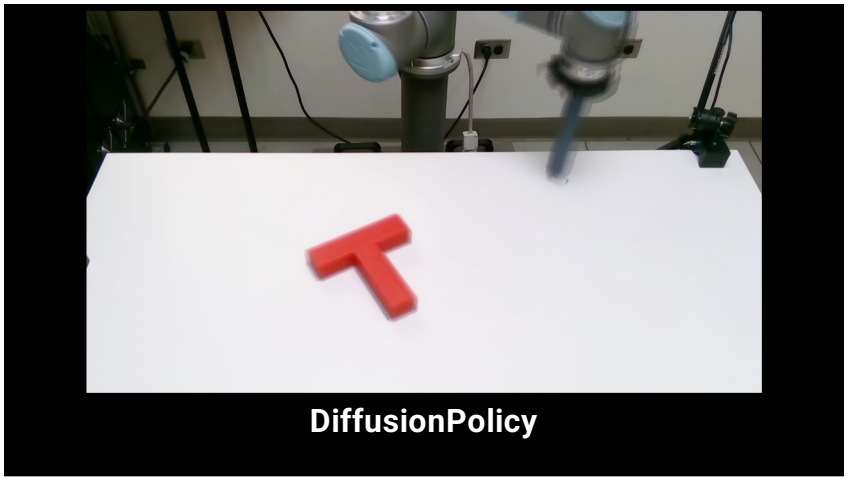

21

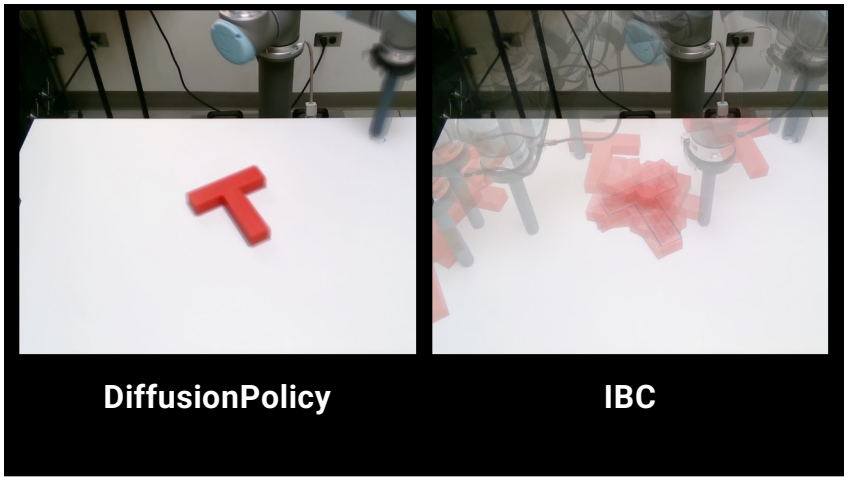

22

Single-Skill Caveat: Robustness is Expensive

23

Behavior Generation with LfD is Real

- **Flexible**
  - If you can teleop it - you can learn it
  - Intuitive to leverage physics
  - Straightforward to add feedback
- **Accessible**
  - Common sense
  - Fast turn around (~1 day)
- **Practical**
  - Mostly works out of the box
  - Robustness, but at a cost
  - Long horizon task

24

Insights

- A **generic** recipe for *decent* but *narrow* skills
- Reliability is still difficult
- Data quality vs quantity vs diversity
  - Data recipe for scaling?
- Benchmarking remains a serious challenge
- Lots of room for theoretical contribution

Scale is a necessity

25

Learning Multitask Models

26

Diffusion Policy

27

Diffusion Policy with Language Conditioning

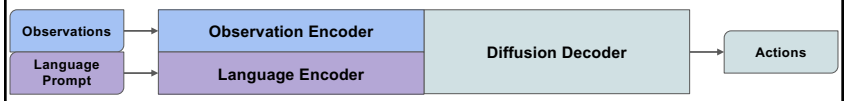

"Stack the bowls."

"Load the dishwasher."

"Flip pancakes."

28

Diffusion Policy with Language Conditioning
Multiple Skills

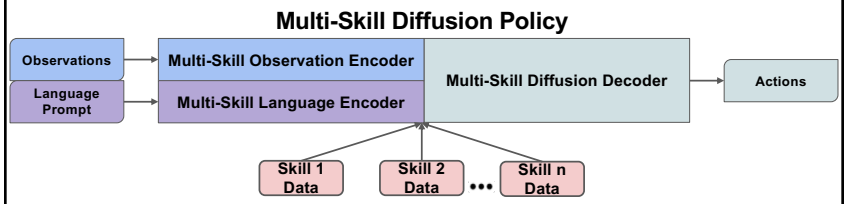

29

Improved Robustness from Multi-Skill Learning

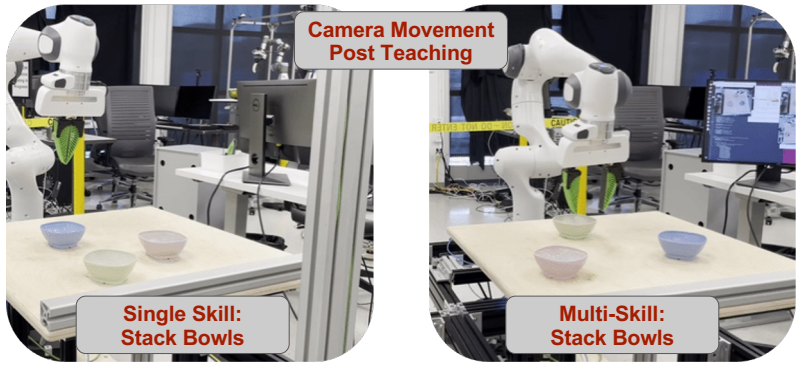

30