Rolling Diffusion Policy

Enhancing Efficiency and Temporal Awareness

Chanhyuk Jung[1]*, Sangwon Kim[2]*, Dasom Ahn[1], In-su Jang[2], Kwang-Ju Kim[2], Sungkeun Yoo[1], Byoung Chul Ko[1]

[1]Keimyung University, South Korea
[2]ETRI - Electronics and Telecommunications Research Institute, South Korea

ICRA IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION ATLANTA 2025

계명대학교 KEIMYUNG UNIVERSITY

ETRI 한국전자통신연구원 Electronics and Telecommunications Research Institute

1. Motivation

Recent advances in diffusion models have enabled high-quality robotic trajectory generation
However, previous diffusion policies suffer from two major limitations:

- **High Computational Cost:** They require hundreds of iterative denoising steps, making real-time inference infeasible
- **Lack of Temporal Awareness:** Standard diffusion policies ignore growing future uncertainty

Our Goal: Develop a fast, temporally-aware diffusion policy for real-time robotic control

2. Key Contributions

- **Rolling Diffusion Mechanism:** Sliding-window denoising for temporally coherent action refinement
- **Temporal Uncertainty Modeling:** Increasing noise for future actions
- **Real-Time Inference:** Single-step action generation, achieving >100X faster inference
- **Superior Performance:** Competitive or improved results on Push-T benchmark

3. Method: Rolling Diffusion Policy (RDP)

Core Idea: Real-time sampling through sliding-window denoising and rolling updates.

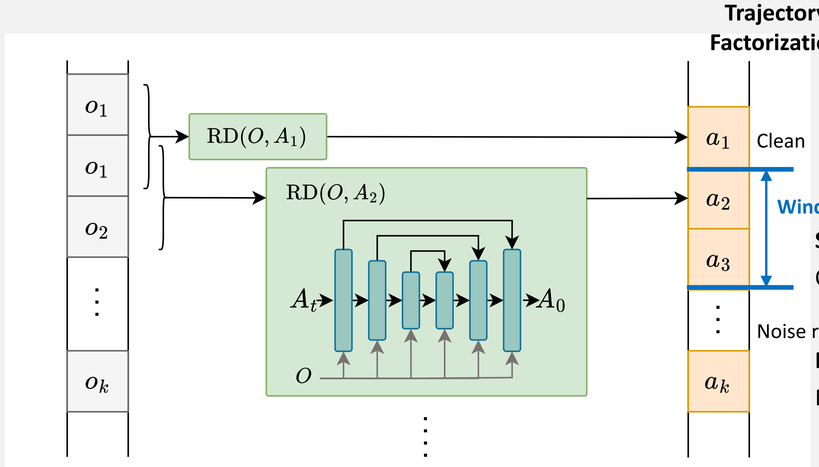

Sliding-Window Denoising:
Only denoise the **active window**; slide forward after each step

Noise Scheduling:
Monotonically increasing noise to model future uncertainty

Result: Single evaluation per action, fast yet high-quality generation.

4. Results

Benchmark: Push-T hybrid (visuomotor control)
Quantitative Results:

Models	Avg. score	Latency (ms)
Diffusion Policy (DP) [1]	0.91	110
Consistency Policy (CP) [7]	0.75	2
Streaming Diffusion Policy (SDP) [5]	0.84	7
Ours	**0.88**	**1**

Observations:

- High action quality (close to DP)
- Drastic latency reduction
- Real-time action generation with single evaluation

5. Visualization

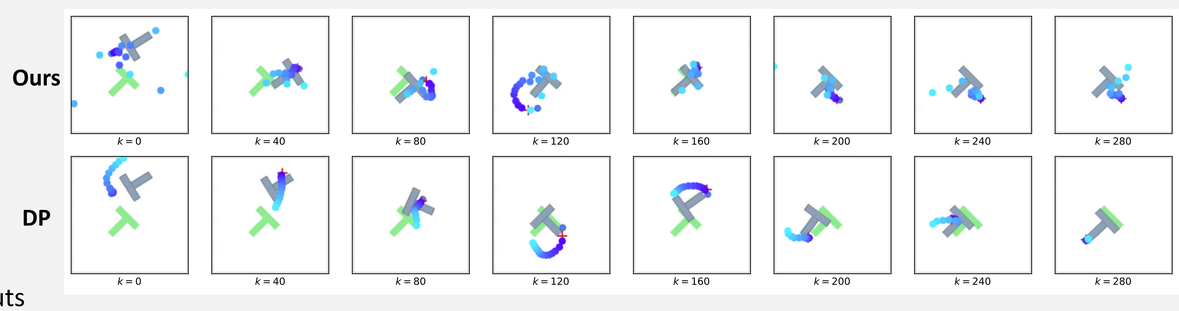

Takeaways:

- Stable long-horizon rollouts
- Computationally efficient sliding-window denoising

6. Conclusion and Future Directions

Conclusion:

Rolling Diffusion Policy (RDP) enables real-time, temporally-aware robotic action generation with high efficiency

Future Directions:

- Dynamic and complex environments
- Stability over long-horizon rollouts
- Integration with hierarchical task planners