Flow Matching Imitation Learning for Multi-Support Manipulation

Quentin Rouxel, Andrea Ferrari, Serena Ivaldi, Jean-Baptiste Mouret

Inria, Université de Lorraine, CNRS, Loria, F-54000.

Project Website

https://hucebot.github.io/flow_multisupport_website/

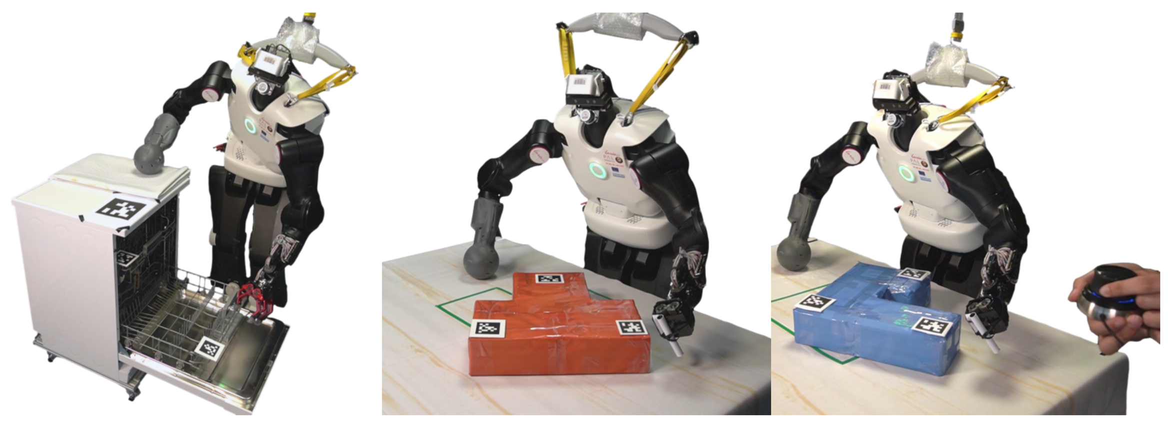

Multi-Support Manipulation

- Use additional contact to **extend manipulation capabilities**
- Reach further, be more stable, apply higher pushing forces
- Leverage whole-body motion and multi-contact strategies

Problem: How to chose contact placement and sequence?

Main Ideas

- **Imitation learning** from human teleoperated demonstrations (behavior cloning)
- Use human "common sense" to learn where and when to place additional contact
- Solve out-of-distribution tasks not covered by demonstrations using shared autonomy assisted teleoperation

Imitation Learning with Flow Matching

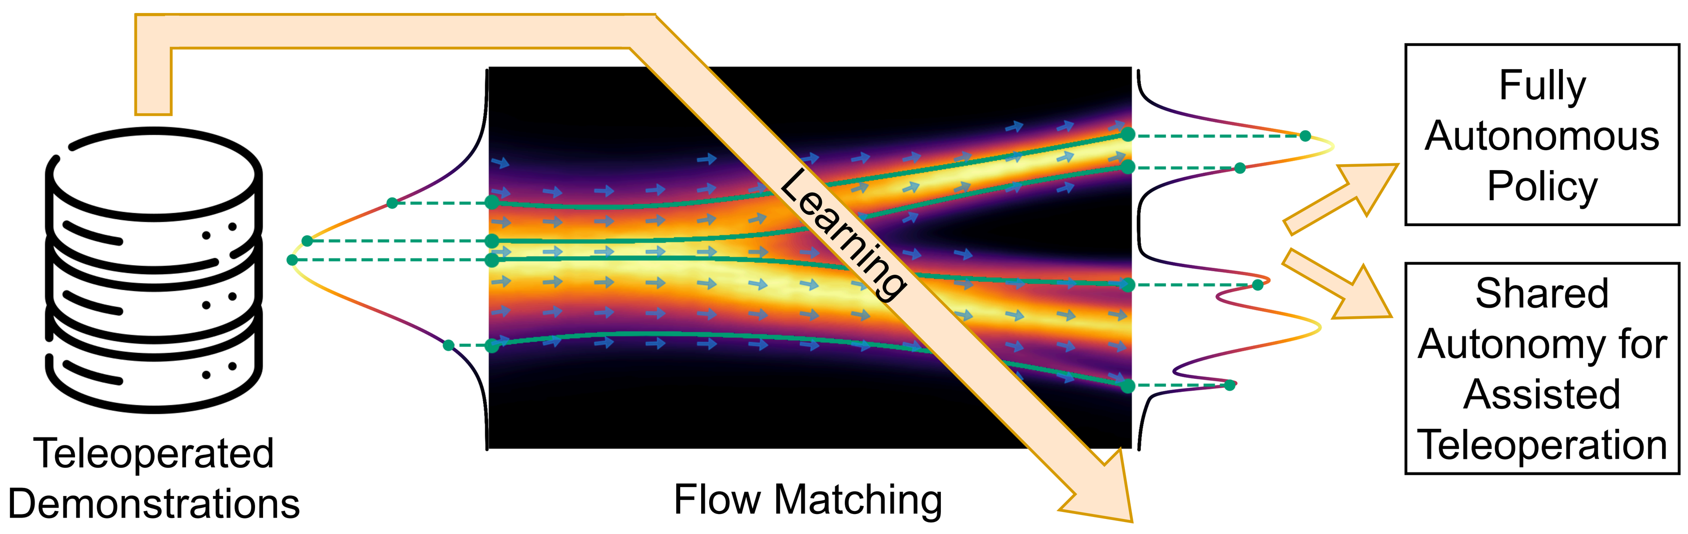

- Use Flow Matching, a generative method able to capture high-dimensional and multi-modal distributions
- Simpler framework, slightly more performant than Diffusion grounded in optimal transport theory
- The policy outputs a full command trajectory instead of a single action

Behavioral Cloning Policy

Policy Inputs Outputs

$$\text{Policy } \pi : \boldsymbol{s}_k \longrightarrow \boldsymbol{a}_k$$

- **Inputs**: effector last pose command, effector contact state, detected marker poses
- **Outputs**: pose commands trajectory, contact trigger trajectory

Training with Flow Matching

$$\boldsymbol{a}^{\text{src}} \sim \mathcal{P}^{\text{src}},\ \boldsymbol{a}^{\text{dst}} \sim \mathcal{P}^{\text{dst}},\ t \sim \mathcal{U}[0,1],$$
$$\boldsymbol{z}_t = (1-t)\boldsymbol{a}^{\text{src}} + t\boldsymbol{a}^{\text{dst}},$$
$$\mathcal{L}_{\text{flow}} = \mathbb{E}_{\boldsymbol{a}^{\text{src}},\boldsymbol{a}^{\text{dst}},t} \left\| f(\boldsymbol{z}_t, t, \boldsymbol{s}) - (\boldsymbol{a}^{\text{dst}} - \boldsymbol{a}^{\text{src}}) \right\|^2$$
$$\text{Flow } f : \boldsymbol{z}, t, \boldsymbol{s} \longrightarrow \Delta \boldsymbol{z},$$

Multi-Contact Low-Level Controller

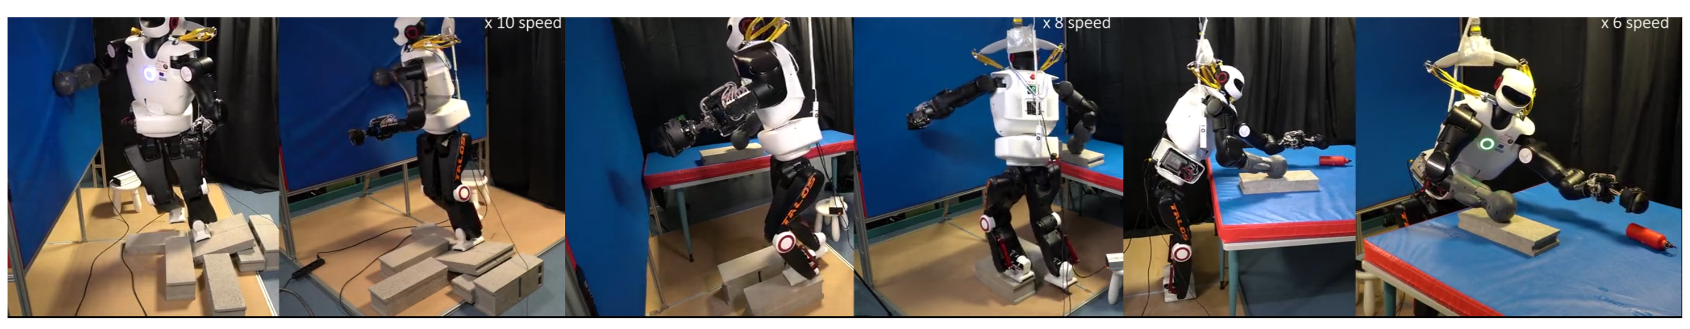

- Use our previous work **SEIKO** Retargeting and Controller [1, 2]
- Retargeting from Cartesian to whole-body commands with feasibility constraints
- Realize multi-contact on position-controlled robot
- Regulate contact forces for smooth contact switches

Experiments on Talos Humanoid Robot

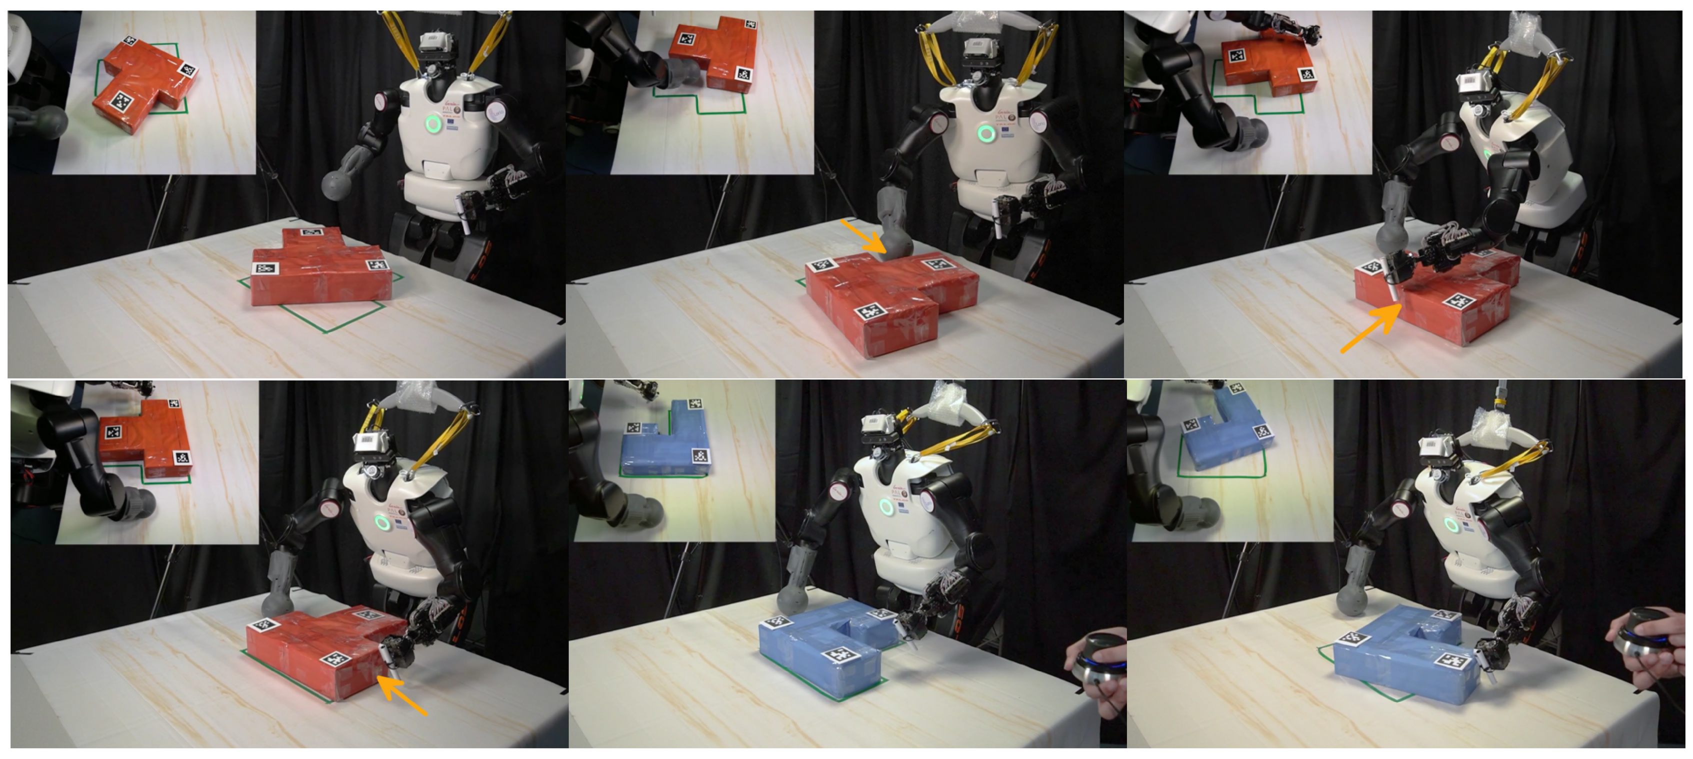

The robot uses both hands for non-prehensile manipulation, making and breaking contact

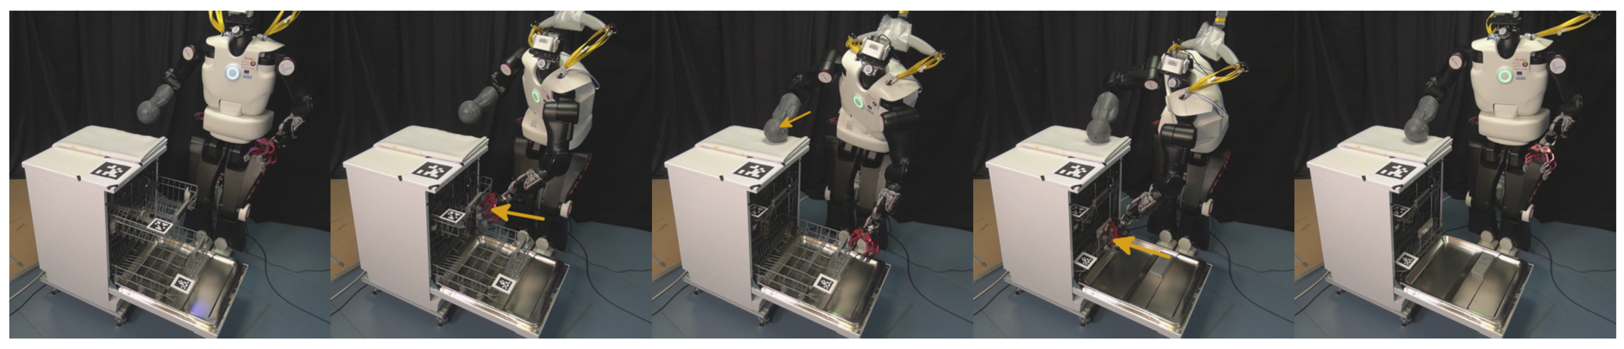

To reach and close the lower drawer, the robot has learned to place its right hand on top of the dishwasher

Comparison Between Flow Matching and Diffusion

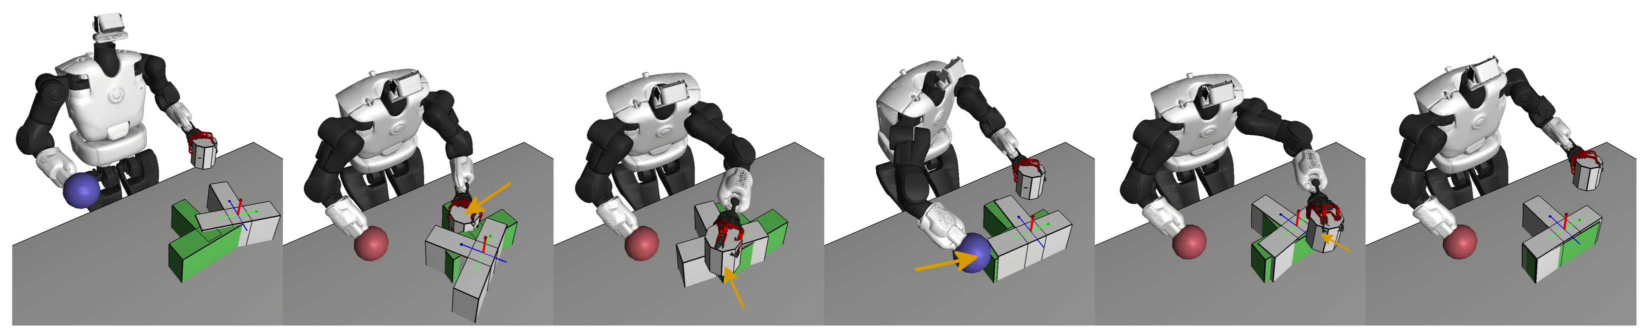

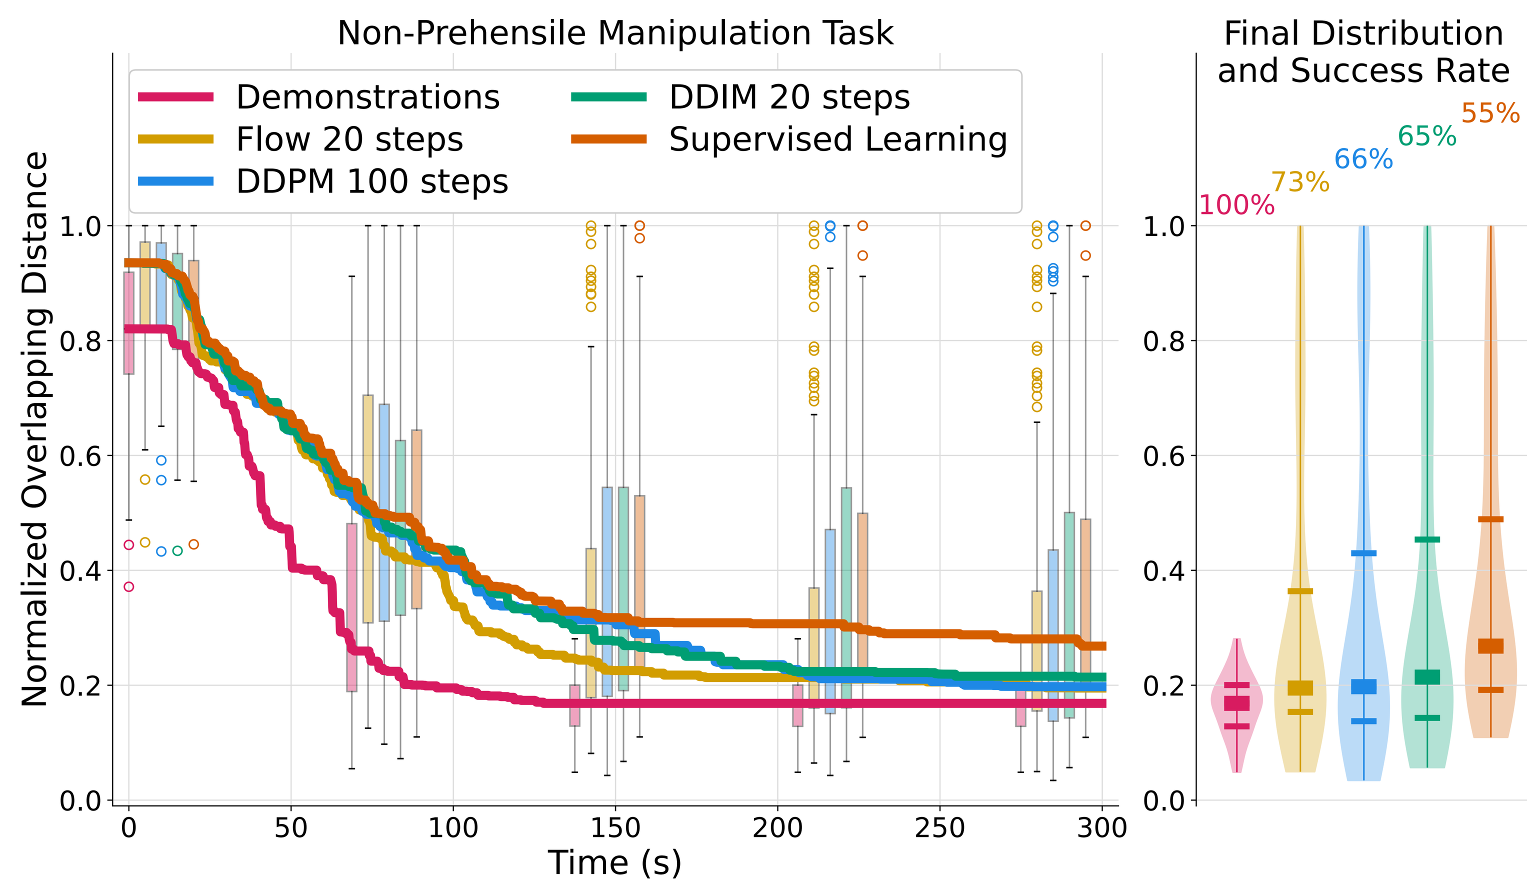

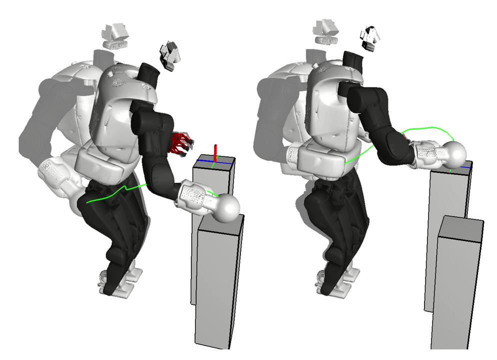

Method	Inference Time (ms)	In Distribution		Out of Distribution	
		Success Rate	Median [Q1, Q3] (cm)	Success Rate	Median [Q1, Q3] (cm)
Demonstrations	–	100%	1.3 [0.9, 2.0]	–	–
Flow 20 steps	35 ± 4	99%	1.4 [1.0, 2.1]	**78%**	**3.4 [1.9, 5.7]**
DDPM 100 steps	178 ± 12	100%	1.5 [0.9, 1.8]	69%	4.0 [2.4, 5.9]
DDIM 20 steps	39 ± 4	100%	1.4 [0.9, 2.0]	67%	3.9 [2.7, 6.1]
Supervised Learning	$\mathbf{3 \pm 1}$	92%	4.1 [2.6, 5.3]	52%	7.6 [4.3, 12.3]

- Both Flow Matching and Diffusion methods outperform the supervised learning behavior cloning baseline
- Flow method tends to marginally outperform its Diffusion counterparts at comparable inference times

References

[1] Q. Rouxel et al. "Multicontact motion retargeting using whole-body optimization of full kinematics and sequential force equilibrium". In: *IEEE/ASME T-MECH* (2022).

[2] Q. Rouxel et al. "Multi-Contact Whole-Body Force Control for Position-Controlled Robots". In: *IEEE RA-L* (2024).