StRiDE: State-space Riemannian Diffusion for Equivariant Planning

Nishanth Rao*, Evangelos Chatzipantazis*, Kostas Daniilidis



Penn Engineering | GRASP Laboratory
General Robotics, Automation, Sensing & Perception Lab

ARCHIMEDES

Diffusion Models for Planning

- Planning formulated as *Trajectory Posterior Sampling*

$$\log p_w(\tau \mid C) = \underbrace{\log p_w(\tau)}_{\substack{\text{diffusion prior} \\ \text{(from expert plans)}}} + \underbrace{\log p(C \mid \tau)}_{\substack{\text{motion optimization} \\ \text{likelihood } \propto \exp(-J(\tau))}} + \text{const.}$$

- implemented via *diffusion guidance.*

$$\text{State} : s = (\theta_1, \ldots, \theta_K, \omega_1, \ldots, \omega_K)$$
$$\text{Trajectory} : \tau := (s_1, \ldots, s_T)$$
$$\tau^N \sim \mathcal{N}(0, I)$$
$$\tau^{i-1} = \mu_w(\tau^i, i) - \nabla_\tau J^{i-1}(\tau = \mu_w(\tau^i, i)) + \tilde{\beta}_i \mathcal{N}(0, I)$$

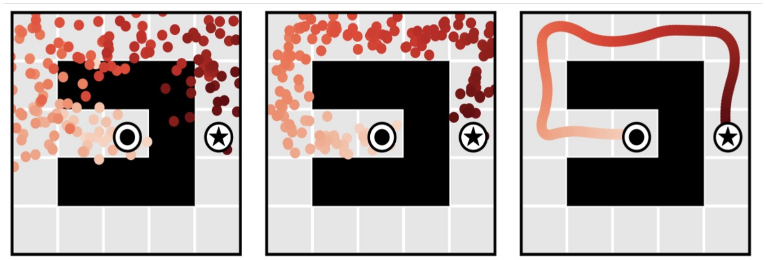

[Janner et al. 2022]

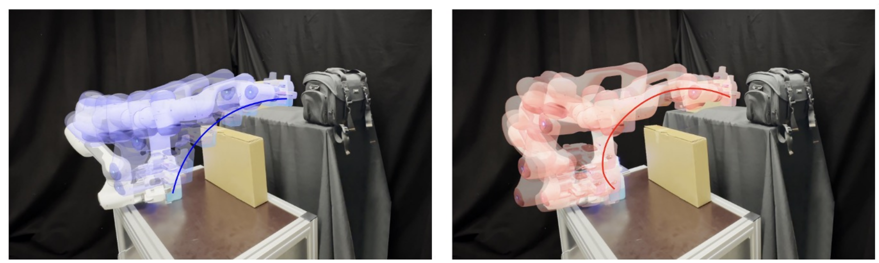

[Joao Carvalho et al. 2024]

Diffusion Models for Planning

$$\text{State}: s = (\theta_1, \dots, \theta_K, \omega_1, \dots, \omega_K)$$
$$\text{Trajectory}: \tau := (s_1, \dots, s_T)$$
$$\tau^N \sim \mathcal{N}(0, I)$$
$$\tau^{i-1} = \mu_w(\tau^i, i) - \nabla_\tau J^{i-1}(\tau = \mu_w(\tau^i, i)) + \tilde{\beta}_i \mathcal{N}(0, I)$$

+ Fast Prior Sampling.
+ Trajectory-level Prediction (avoids error accumulation)

- Slow Motion Optimization.
- Not Generalizable to new environments.

❑ Observations:
- Diffusion Guidance dominates the planning time.
  But limits the complexity of the costs
- Generalization to new environments is also taken care of via guidance.
  But limits generalization to small perturbations of an initial environment.

We can solve both issues by addressing the redundancy in the representation in all stages.

Our Proposal

1. Posterior Sampling via State-Space Diffusion.
   - ✓ Respect state-space topology during diffusion, denoising, guidance.
   - ✓ Faster Training: Important bottleneck; retrain for every new experiment.
   - ✓ More effective training: State-space
   - ✓ Connection between Lie-algebra and state space permits a hybrid approach.
   - ✓ Faster Motion Optimization using Riemannian Guidance; major bottleneck.

2. Generalization by Equivariant Planning in Canonicalized Environment.
   - Train on canonical environment and generalize to transformed environments.
   - Here we deal with global SO(2) rotations around the base.
   - Diffusion guidance is still used for local changes.
   - Much faster i.e. fewer guidance steps when local and global combined.

Equivariance + State-Space Diffusion and Guidance

a)

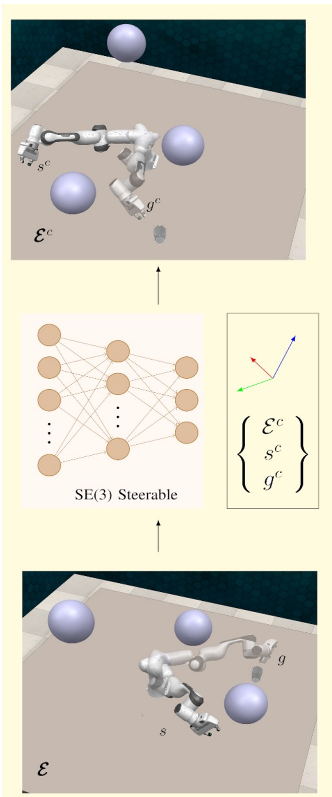

b)

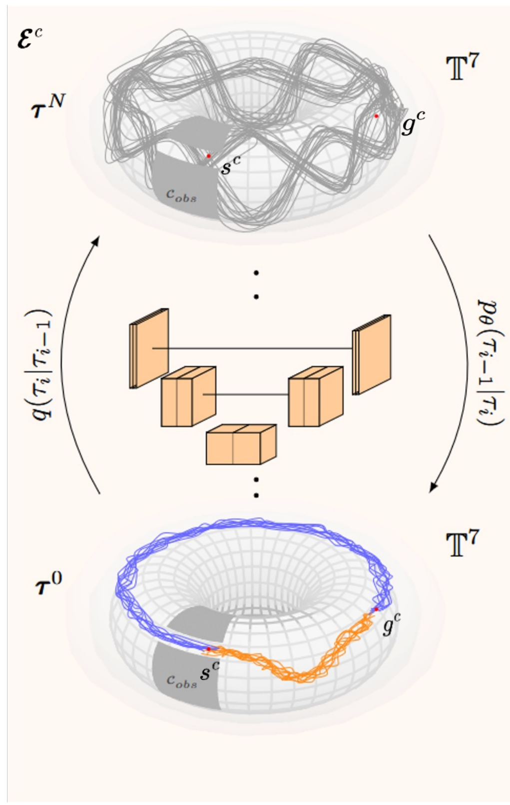

c)

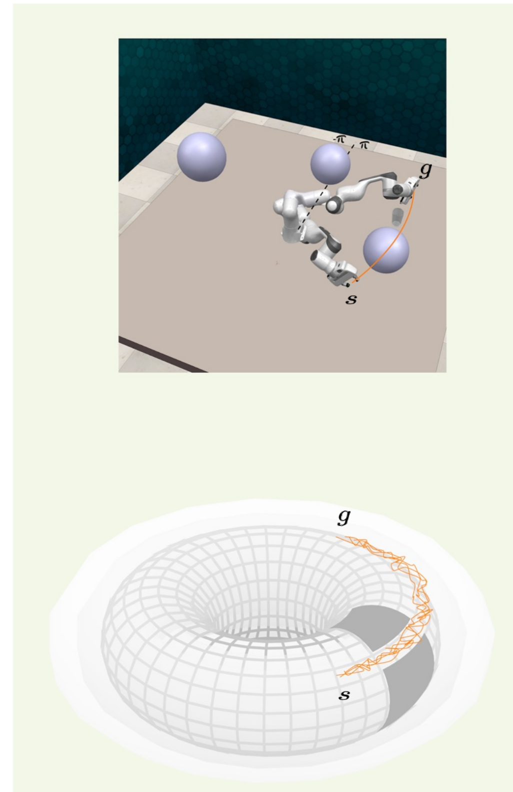

Plan in the canonicalized environment obtained via equivariant frame prediction.

Trajectory Sampling via State-space diffusion

De-canonicalize and discard infeasible branch.

StRiDE: State-space Riemannian Diffusion for Equivariant Planning

2 revolute DoF Configuration Space

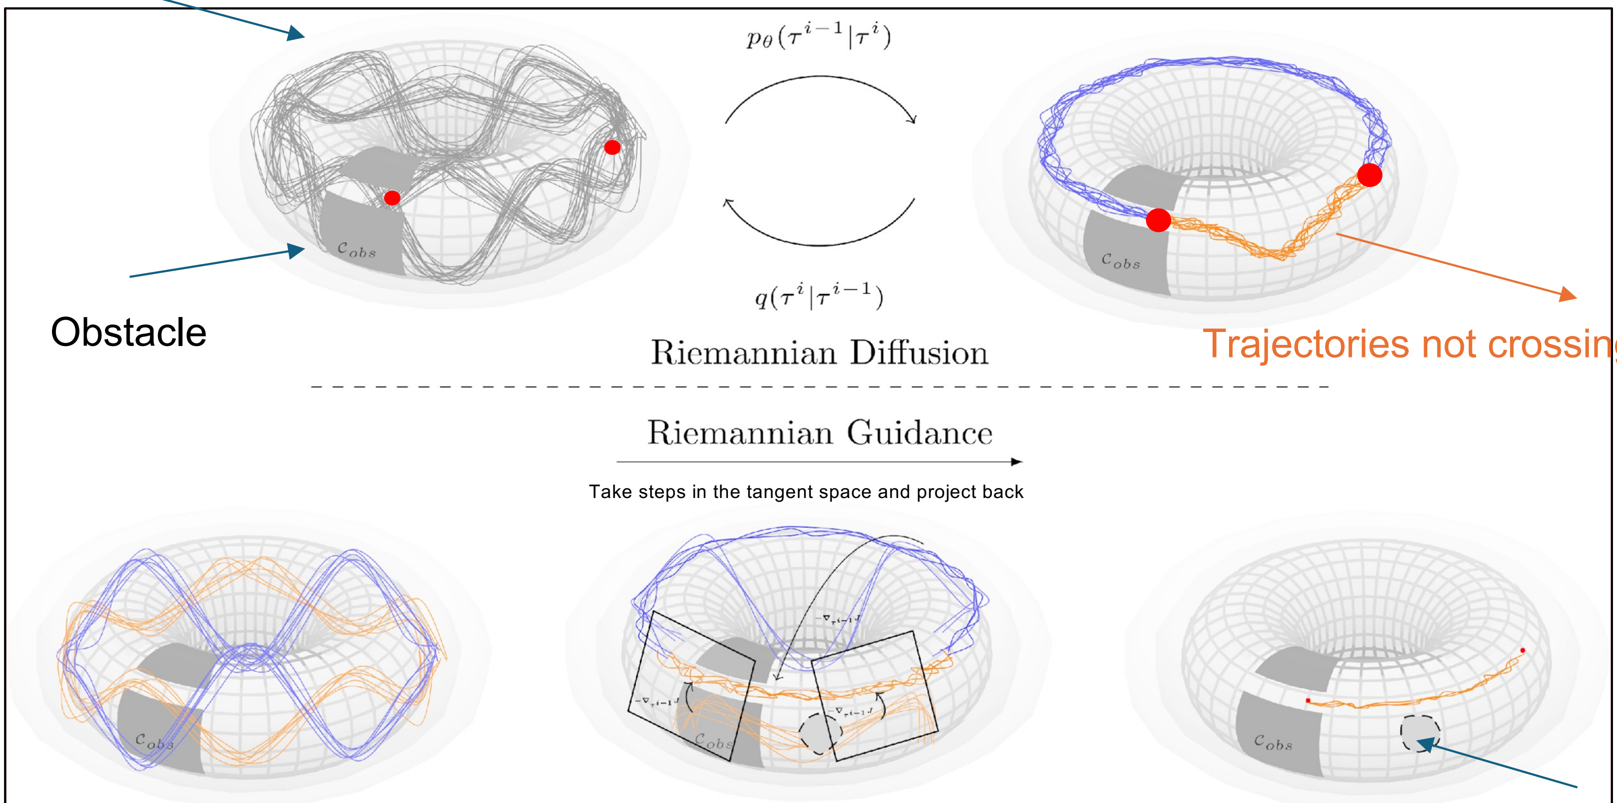

Obstacle

Trajectories not crossing joint limits

New obstacle being avoided by guidance

By exploiting the topology of the state-space manifold we move faster towards feasible paths.

Isotropic Gaussian on SO(2) + Wrapped Normal Distribution

- SO(2), the group of 2x2 rotation matrices, is 1d embedded in 4d.

- Probability measure on SO(2):
  - Cannot define via density (does not exist!)
  - Can still define via the sampling (pushforward measure)
- Isotropic Gaussian on SO(2)

$$\int_{X_2} g\, d(f_*\mu) = \int_{X_1} g \circ f\, d\mu.$$

$$\boxed{R \sim \mathcal{IG}_{SO(2)}(R_\mu, \sigma^2) \iff \theta \sim \mathcal{N}(\mu, \sigma^2), R = \exp_m(\theta)}$$

- In the figure, we see the "density" with red constrained on the unit vectors (isomorphic to SO(2)).

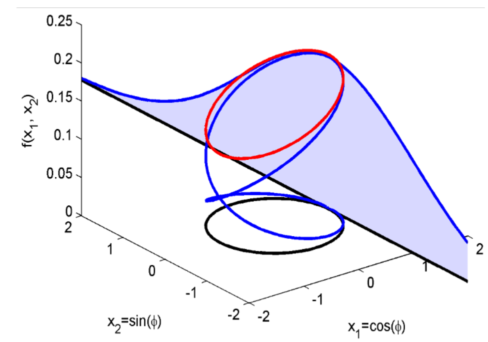

Isotropic Gaussian on SO(2) + Wrapped Normal Distribution

- **Properties of our distribution:**
  - Expectation**:** $R \sim \mathcal{IG}_{SO(2)}(R_\mu, \sigma^2) \implies \mathbb{E}[R] \propto R_\mu$
  - Closure under SO(2) convolution: $R_{1/2} \sim \mathcal{IG}_{SO(2)}(R_{\mu_{1/2}}, \sigma^2_{1/2}) \implies R_1 R_2 \sim \mathcal{IG}_{SO(2)}(R_{\mu_1} R_{\mu_2}, \sigma_1^2 + \sigma_2^2)$

- Connection with directional statistics. Density, if we return back to angles via log map?

$$R \sim \mathcal{IG}_{SO(2)}(R_\mu, \sigma^2) \implies \log R \sim \mathcal{WN}(\mu \ mod 2\pi, \sigma^2)$$

- Wrapped Normal:

$$\boxed{\mathcal{WN}(\theta; \mu, \sigma^2) = \sum_{k=-\infty}^{\infty} \mathcal{N}(\mu + 2\pi k, \sigma^2), \ \ \theta, \mu \in [-\pi, \pi)}$$

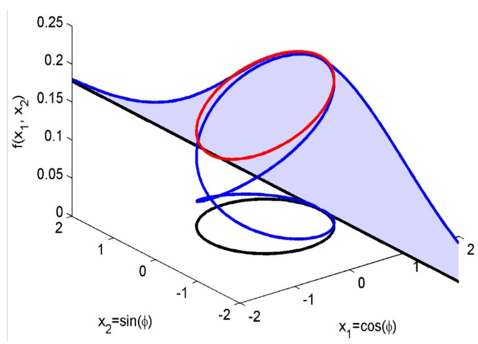

StRiDE: Diffusion Stage

- State-space Representation: Use intrinsic parameters but embed in Euclidean.

$$\boxed{\begin{array}{l}\text{State} : s = (\cos\theta_1, \sin\theta_1, \ldots, \cos\theta_K, \sin\theta_K, \omega_1, \ldots, \omega_K) \\ \text{Trajectory} : \tau := (s_1 = s_{start}, \ldots, s_T = s_{goal})\end{array}}$$

$$\boxed{\mathcal{D} = \{(\tau_\theta^0, \tau_\omega^0)\}}$$

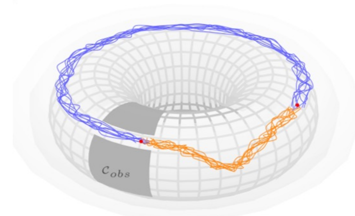

- Isotropic Gaussian for on-manifold diffusion

$$\boxed{R \sim \mathcal{IG}_{SO(2)}(R_\mu, \sigma^2) \iff \theta \sim \mathcal{N}(\mu, \sigma^2), R = \exp_m(\theta)}$$

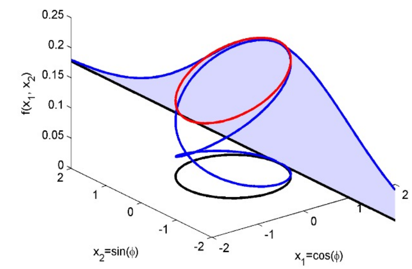

- Single-step Multi-Scale on-manifold Diffusion:

$$\boxed{\begin{array}{ll} R^{\epsilon_i} \sim \mathcal{IG}^{KT}_{\mathrm{SO}(2)}(I, (1-\bar{\alpha}_i)), & \Omega^{\epsilon_i} \sim \mathcal{N}(0, (1-\bar{\alpha}_i) I) \\ \tau_\theta^i = R^{\epsilon_i} \operatorname{Exp}(\sqrt{\bar{\alpha}_i}\, \tau_\theta^0), & \tau_\omega^i = \sqrt{\bar{\alpha}_i}\, \tau_\omega^0 + \Omega^{\epsilon_i}. \end{array}}$$

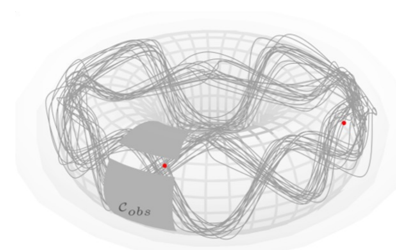

StRiDE: Data Collection and Training

- Data from (expert) RRT-Connect :

$$\mathcal{D} = \{(\tau_\theta^0, \tau_\omega^0)\} \sim P(s_{start}, s_{goal}, \mathcal{E}, T)$$

- Collect trajectories. Both:
  - positive (feasible for canonical environment) and
  - negative (feasible for some rotated environment)

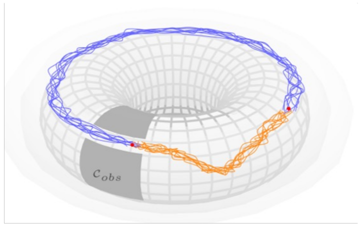

- Training Loss-distance on state-space manifold:

$$\mathcal{L}(w) = \mathbb{E}\Big[d\Big(\mathrm{Exp}\Big(\frac{1}{\sqrt{1-\bar{\alpha}_i}} \log R^{\epsilon_i}\Big), \epsilon_w^\theta(\tau^i, i)\Big) + \big\|\Omega^{\epsilon_i} - \epsilon_w^\omega(\tau^i, i)\big\|_2^2\Big].$$

StRiDE: Architecture

- Input:
  Key: Input and output trajectory representation:
    Continuous Trajectories stay continuous.
    There is no redundancy.
    Trajectories processed as vectors conforming to standard pipelines.

- Output:
  - Embedded Trajectory Noise $\boxed{\begin{aligned}\epsilon_j &= (\cos\theta_1^\epsilon, \sin\theta_1^\epsilon, \ldots, \sin\theta_n^\epsilon, \omega_1^\epsilon, \ldots, \omega_n^\epsilon)\\ \epsilon_\tau &= (\epsilon_1, \ldots, \epsilon_T)\end{aligned}}$
  - Can Normalize network output; implicitly done when we return to lie algebra.

- Architecture (U-Net like):
  - 1D-Conv across time: Locality helps Compositionality.
  - Fully Connected across states.

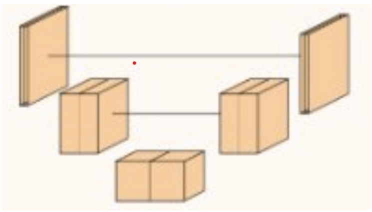

StRiDE: Denoising and Guidance

- On-manifold Denoising:

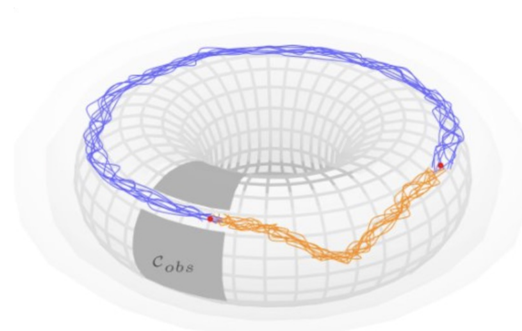

$$
{}^{R}\mu_w({}^{R}\tau^i, i) = \exp\left(\frac{1}{\sqrt{\bar{\alpha}_i}}\log {}^{R}\tau^i\right)\exp\left(-\frac{1}{\sqrt{1-\bar{\alpha}_i}} Log\epsilon_w^q(\tilde{\tau}^i, i)\right)
$$

$$
R_\mu({}^{R}\tau^i, i; w) = \exp\left(\frac{\sqrt{\alpha_{i-1}}(1-\bar{\alpha}_{i-1})}{1-\bar{\alpha}_i}\log {}^{R}\tau_\theta^i\right)\exp\left(\frac{\sqrt{\bar{\alpha}_{i-1}}\beta_i}{1-\bar{\alpha}_i}\log {}^{R}\mu_w({}^{R}\tau^i, i)\right)
$$

- On-manifold guidance via Riemannian Gradient Descent:

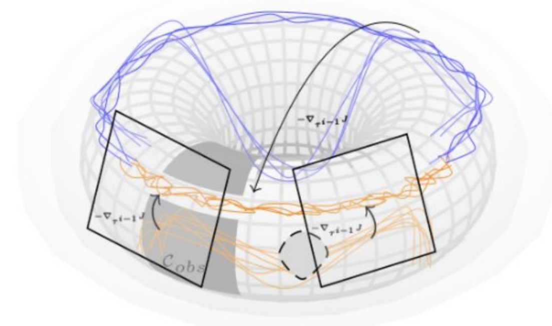

Interleave Denoising with Guidance
Gradient step on state-space:

$$
\tau_\theta^{i-1} = \tau_\theta^i \exp\left(-\eta_i \nabla_\lambda J(\tau_\theta^i \exp(\lambda), \tau_\omega^i)|_{\lambda=0}\right)
$$

$$
J : SO(2)^{nT} \times \mathbb{R}^{nT} \to \mathbb{R}_+
$$

Equivariance by canonicalization

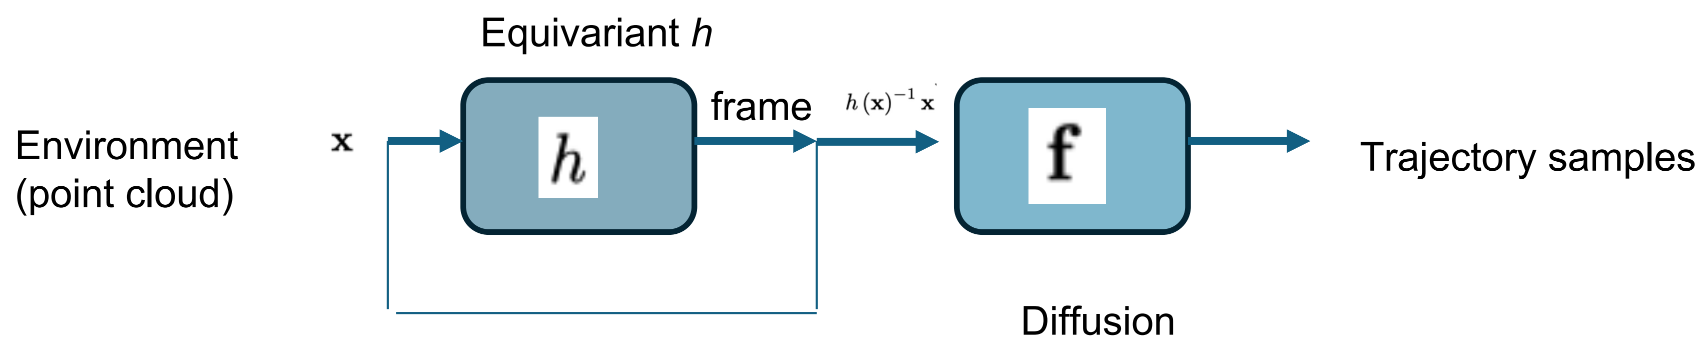

$$\phi(\mathbf{x}) = h'(\mathbf{x})\,\mathbf{f}\left(h(\mathbf{x})^{-1}\mathbf{x}\right)$$

$$h'(\mathbf{x}) = \rho'\left(\rho^{-1}\left(h(\mathbf{x})\right)\right)$$

Equivariance with Learned Canonicalization Functions

Sékou-Oumar Kaba [*1,2] **Arnab Kumar Mondal** [*1,2] **Yan Zhang** [3] **Yoshua Bengio** [4,2] **Siamak Ravanbakhsh** [1,2]

Our Proposal 2: Generalization via Equivariance

- Plan in **canonical environment** and transform the plan **back to the original**.
- Here we deal with **global** SO(2) rotations around the base.
  - Symmetry breaking due to angle limits **(positive and negative trajectories)**
- Fewer guidance steps when **local** and **global** combined.

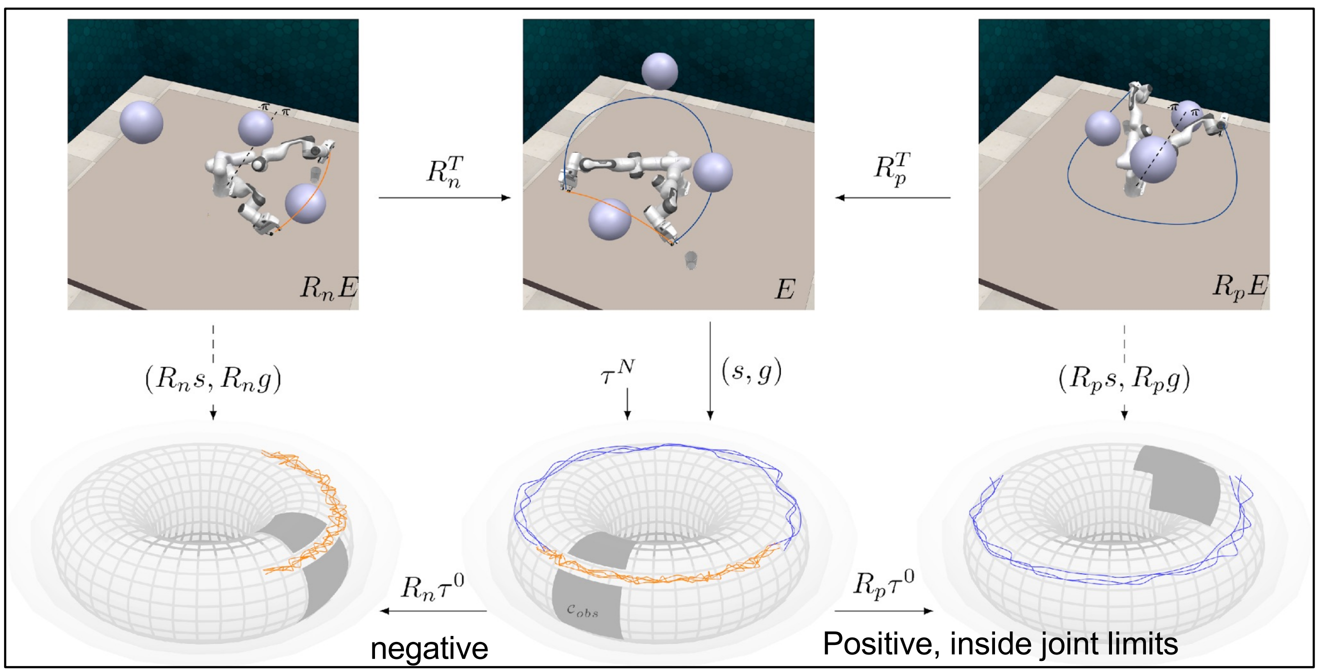

Pandas Spheres Environment:

First row depicts results in the canonical environment.
The second row indicates results averaged over transformed environments

	$\mathcal{S} \uparrow$	$\mathcal{C}_s \downarrow$	$\mathcal{C}_p \downarrow$	$\mathcal{C}_b \downarrow$	$t \downarrow$	
RRTC+GPMP	1.0	–	8.1 ± 1.1	–	226.14 ± 13.4	$\mathbb{E}_{q_i,q_f}[.]$
MPD Canonical	1.0 ± 0.0	7.6 ± 3.41	6.5 ± 2.74	11.54 ± 6.07	23.12 ± 1.1	
STRiDE	1.0 ± 0.0	8.7 ± 1.7	7 ± 1.45	12.6 ± 2.92	9.98 ± 0.9	
MPD Canonical	0.43 ± 0.22	8.77 ± 2.82	4.65 ± 1.49	11.15 ± 3.68	–	$\mathbb{E}_{q_i,q_f}\mathbb{E}_g[.]$
STRiDE	0.97 ± 0.03	8.61 ± 1.57	7.01 ± 0.9	12.9 ± 2.31	–	
EQ-Prior-Guidance	0.85 ± 0.17	9.3 ± 1.46	7.51 ± 1.83	13.96 ± 3.02	10.31 ± 1.85	

Table 1: Metrics are reported as $(\mu \pm \sigma)$ with $\uparrow$ indicate that higher values are better, and with $\downarrow$ indicate that lower values are better.

- $\mathcal{S}$ - success rates, $\mathcal{C}_s$ – smoothness cost, $\mathcal{C}_p$ – path cost, $\mathcal{C}_b$ - overall best cost, t – planning time

- Ablation: We converge to higher success rates faster (with fewer guidance steps) due to Riemannian guidance

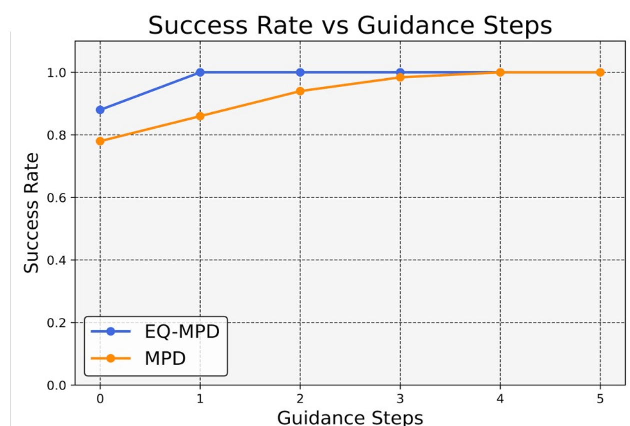

Hardware Experiments

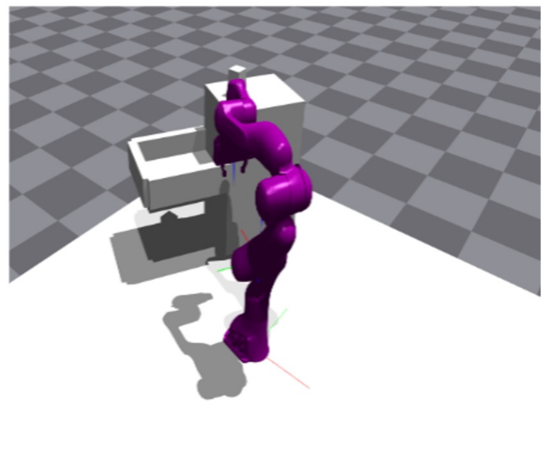

- Canonical Environment in Simulation
- Random rotation of 148.6^{o}
- Objects modelled as differentiable SDFs

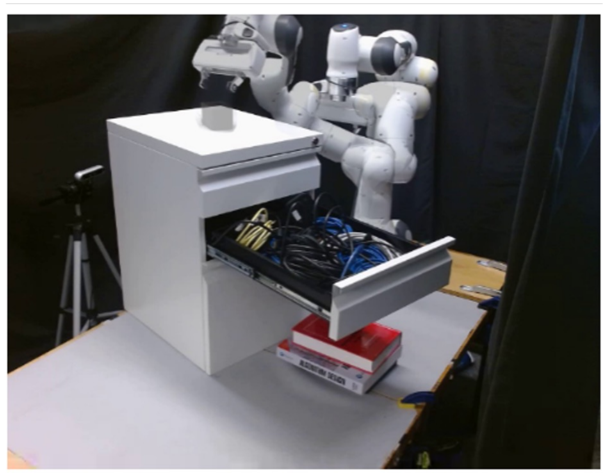

- Real-world experiment, with random (global) rotation of 17.2^{o}
- Pick object under drawer and place on top

- Real-world experiment, with random (global) rotation of 17.2^{o}, with obstruction on top (local transformation)
- Pick object under drawer and place inside

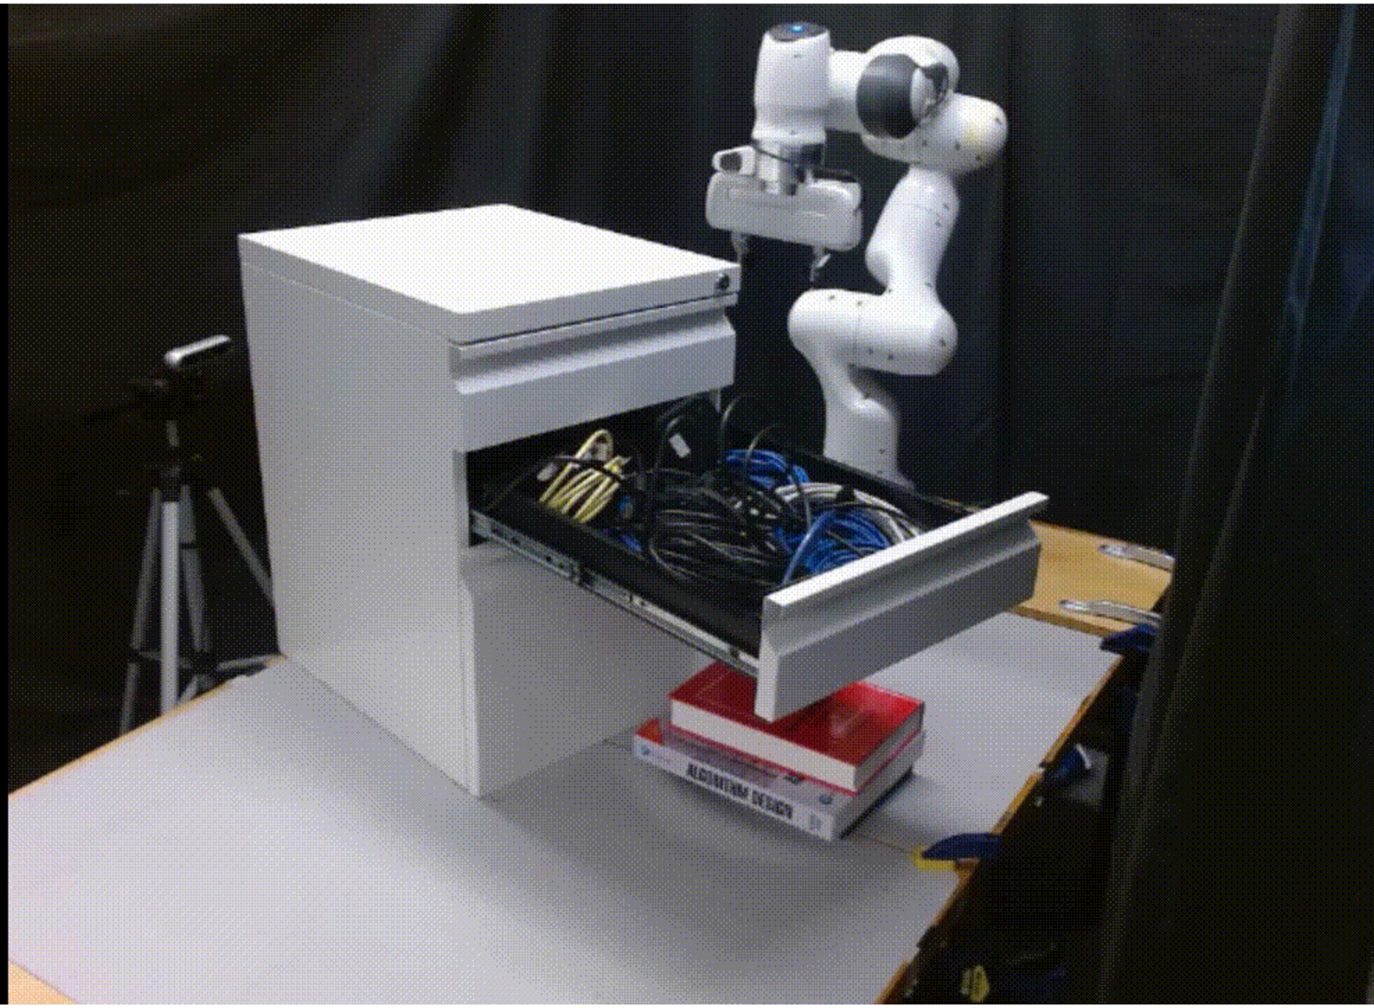
ALGORITHM DESIGN

New environment without retraining

Take home: Canonicalization + Riemannian Diffusion + Riemannian Guidance

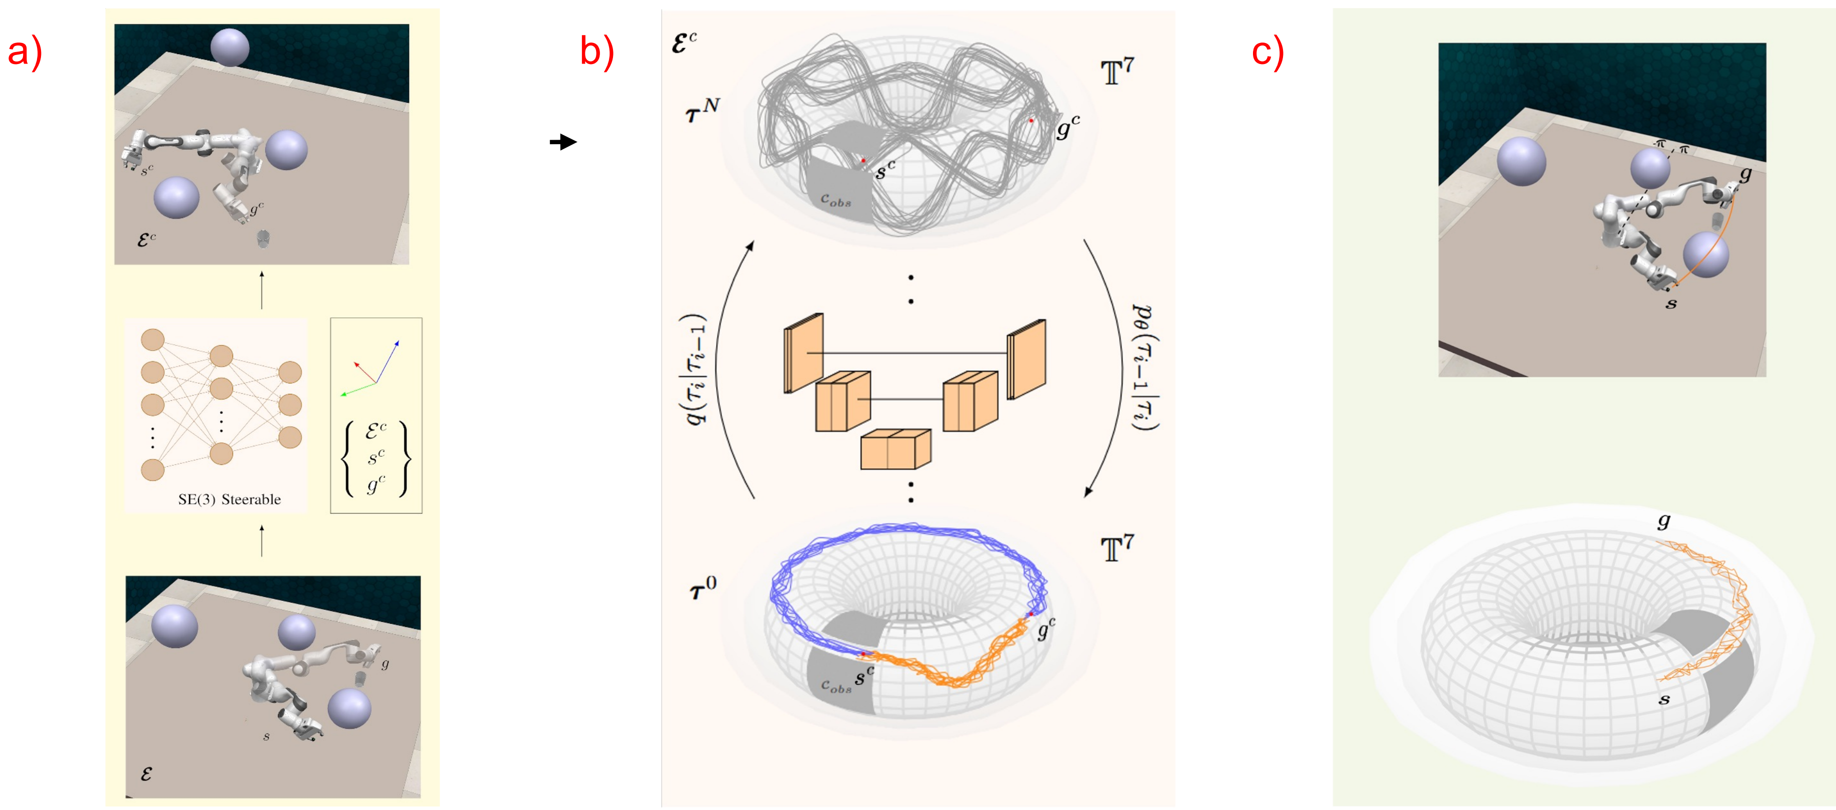

Plan in the canonicalized environment obtained via equivariant frame prediction.

Trajectory Sampling via State-space diffusion

De-canonicalize and discard infeasible branch.