# PRIMP: PRobabilistically-Informed Motion Primitives for Efficient Affordance Learning from Demonstration

Sipu Ruan*, Weixiao Liu†, Xiaoli Wang*, Xin Meng* and Gregory S. Chirikjian*‡
* Department of Mechanical Engineering, National University of Singapore, Singapore
† Deparement of Mechanical Engineering, Johns Hopkins University, Baltimore, MD, USA
‡ Corresponding author, Email: mpegre@nus.edu.sg

***Abstract*—This paper proposes a learning-from-demonstration method using probability densities on the workspaces of robot manipulators. The method, named "PRobabilistically-Informed Motion Primitives (PRIMP)", learns the probability distribution of the end effector trajectories in the 6D workspace that includes both positions and orientations. It is able to adapt to new situations such as novel via-point poses with uncertainty and a change of viewing frame. The method itself is robot-agnostic, in which the learned distribution can be transferred to another robot with the adaptation to its workspace density. The learned trajectory distribution is then used to guide an optimization-based motion planning algorithm to further help the robot avoid novel obstacles that are unseen during the demonstration process. The proposed methods are evaluated by several sets of benchmark experiments. PRIMP runs more than 5 times faster than the compared existing probabilistic methods while generalizing trajectories more than twice as close to both the demonstrations and novel desired poses. It is then combined with our robot imagination method that learns object affordances, illustrating the applicability of PRIMP to learn tool use through physical experiments.**

## I. Introduction

For a robot to be truly intelligent, it needs the ability to learn from prior knowledge while adapting to unseen scenarios. The prior knowledge can be from human-demonstrated motions, which are difficult to pre-program but are ubiquitous in household environments, like scooping powder (as in Fig. 1). It is related to a popular field in robot learning, namely *Learning-from-Demonstration (LfD)* [21] or *Programming-by-Demonstration (PbD)* [4]. Many works on LfD encode demonstrations as trajectories in Euclidean space [16, 30]. However, the models in the full workspace (including both position and orientation) of a robot manipulator have not been considered until recently [29, 5, 14, 23, 2]. Our work focuses on the robot workspace and proposes a novel method using probability densities on Lie groups, denoted as *PRobabilistically-Informed Motion Primitives (PRIMP)*. The mathematical model is inspired by a concept initially introduced in our group more than 25 years ago, the concept of loop entropy [8, 6] and more recent work on inverse reachability mapping [24, 15]. It is further extended here into LfD with via-point conditioning as compared to only subjecting to end constraints. The learned knowledge is robot-agnostic but can be transferred among different robots. It is also able to deal with extrapolation cases

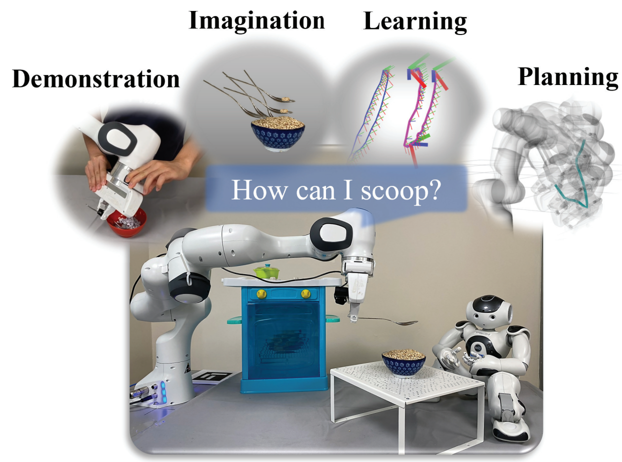


Fig. 1: The robot arm is asked to use a spoon to scoop from a bowl in a household environment. With the help of human demonstrations, imagination of object affordance, learning skills from the demonstrations and motion planning, the robot is able to fulfill the task in a novel scene with unseen obstacles.

and model with a few or even a single demonstration.

When there are some novel obstacles unseen during the demonstrations, only using LfD methods is not enough to generate feasible trajectories. Guided motion planning, which combines LfD and motion planning, has become popular in the recent decade [18, 28, 19, 11]. The goal is to make the motion collision-free while keeping the critical features from the learned trajectory as many as possible. In our work, an optimization-based planner, *Stochastic Trajectory Optimization for Motion Planning (STOMP)* [17], is applied as the base framework. The learned trajectory distribution via PRIMP is used as a reference. A novel cost function with respect to this reference distribution is proposed to guide the planning process. Instead of joint space, the cost function is based on the workspace of the robot end effector. Therefore, the planner is named as *Workspace-STOMP*.

Apart from novel obstacles, the object that the robot interacts with might also be unseen. For example, the robot has been demonstrated how to pour powders from a cup into a bowl. In this case, the cup is treated as a tool for pouring and the bowl has the affordance of containing. In a new scenario, the tool might be changed into a spoon and the object becomes a vase, which has the same affordance of containing. The learned trajectory distribution should be able to adapt to this new situation with the same set of demonstrations, even when the appearances and categories of the tool and object are totally different. In order to fulfill similar tasks intelligently, the understanding of object functionality and affordance is a key aspect [13, 27, 25]. Our work learns the key points of a task via physics simulation, which is used as new goal or via-point poses for re-production. The *via-point pose* is defined as a pose at an intermediate time step that the robot is desired to pass through. The affordance learning of an object is then combined into a robotic system with PRIMP and Workspace-STOMP through physical experiments.

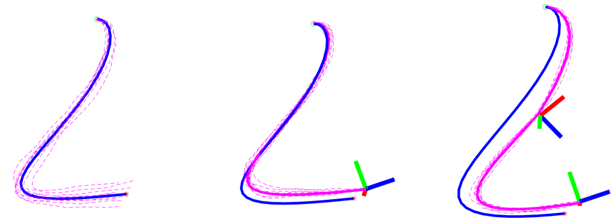

(a) Original trajectory mean and samples (b) Condition on goal pose (c) Condition on goal and a via-point pose

Fig. 2: Examples of the conditional probability on new via-point poses with uncertainty. The solid blue and magenta curves are the means of the encoded joint and conditioned distribution, respectively. Dashed magenta curves are the random trajectory samples from the probability distribution.

## II. PROBABILISTICALLY-INFORMED MOTION PRIMITIVES

The trajectory is represented discretely by a user-selected number of time steps (*i.e.*, $N_{\text{step}}$). Each pose is modeled as an element in Lie group $G$. Therefore, the full state is considered in a product space $G \times ... \times G$. The set of demonstrated trajectories is denoted as $\left\{ g_i^{(k)} \in \text{SE}(3) \right\}$, where $i$ is the step number and $k$ is the index of the demonstration. The process is assumed as Markovian, where $i^{\text{th}}$ step is only affected by its neighboring $(i-1)^{\text{th}}$ and $(i+1)^{\text{th}}$ steps.

### A. General framework of PRIMP

The probability distribution of the $i^{\text{th}}$ pose with respect to the $(i-1)$-th pose is approximated using a Lie-theoretic method (Sec. II-B). The computed initial mean and covariance is then encoded into the joint distribution of the whole trajectory (Sec. II-C). Several types of adaptations to novel scenarios are then introduced (Sec. II-D): (1) via-point conditioning (Sec. II-D1); (2) equivariance to new viewing points (Sec. II-D2); and (3) fusion with robot-specific workspace density (Sec. II-D3).

### B. Computation of relative pose distribution

For a set of poses in SE(3), the sample mean $\mu_i \in \text{SE}(3)$ satisfies $\sum_{k=1}^m \log(\mu_i^{-1} g_i^{(k)}) = \mathbb{O}$, which can be iteratively solved [1]. The mean trajectory can be directly computed from the demonstrations. The initial covariance $\Sigma_{i,i+1}$ encodes the uncertainty of $(i+1)$-th step given the $i^{\text{th}}$ step. It is estimated by the set of relative poses, *i.e.*, $\left\{ \Delta_{i,i+1}^{(k)} = \left( g_i^{(k)} \right)^{-1} g_{i+1}^{(k)} \right\}$. With this set, the sample covariance can be computed as $\Sigma_{i,i+1} = \frac{1}{m} \sum_{k=1}^m \log^\vee \left( \mu_{i,i+1}^{-1} \Delta_{i,i+1}^{(k)} \right) \log^{\vee T} \left( \mu_{i,i+1}^{-1} \Delta_{i,i+1}^{(k)} \right)$, where the $\vee$ operator extracts the Lie algebra coefficients into a vector (as defined in [7]).

### C. Probabilistic encoding of joint distributions

After computing the trajectory distribution with mean $\{\mu_i \in \text{SE}(3)\}$ and covariance between adjacent steps $\{\Sigma_{i,i+1} \in \mathbb{R}^{6\times 6}\}$ from Sec. II-B, the joint distributions of the whole trajectory can be computed. The key assumption is that the variation of $i^{\text{th}}$ pose only depends on its two neighboring poses and $g_0 = \mu_0$ is fixed. The probability density for Gaussian distributions with small variations can be explicit expressed with the joint variable $\boldsymbol{x}_{1,...,n} \doteq [\boldsymbol{x}_1^T, ..., \boldsymbol{x}_i^T, ..., \boldsymbol{x}_n^T]^T$, where $\boldsymbol{x}_i = \log^\vee(\mu_i^{-1} g_i)$ and joint covariance $\Sigma_{1,...,n}'^{-1}$, where the non-zero elements are

$$
\begin{aligned}
\Sigma_{1,...,n}'^{-1}(i,i) &= \begin{cases} \Sigma_{i-1,i}^{-1} + \widetilde{\Sigma}_{i,i+1}^{-1} & (i = \{1,...,n-1\}) \\ \Sigma_{i-1,i}^{-1} & (i = n) \end{cases} \\
\Sigma_{1,...,n}'^{-1}(i,i+1) &= -Ad_{i,i+1}^{-T} \Sigma_{i,i+1}^{-1} \ (i = \{1,...,n-1\}) \\
\Sigma_{1,...,n}'^{-1}(i+1,i) &= -\Sigma_{i,i+1}^{-1} Ad_{i,i+1}^{-1} \ (i = \{1,...,n-1\}) ,
\end{aligned} \tag{1}
$$

where $\Sigma_{1,...,n}' \in \mathbb{R}^{6n\times 6n}$, $Ad(g)$ is the adjoint operator for $g$ in a Lie group, which is defined as $Ad(g)\widehat{\boldsymbol{x}} \doteq g\widehat{\boldsymbol{x}}g^{-1}$, $\widehat{\cdot}$ is the inverse operation of $\vee$ for elements in Lie algebra and $Ad_{i,i+1} \doteq Ad(\mu_i^{-1}\mu_{i+1})$ is the adjoint operator for the relative poses between $\mu_i$ and $\mu_{i+1}$ and $\widetilde{\Sigma}_{i,i+1} = Ad_{i,i+1}\Sigma_{i,i+1}Ad_{i,i+1}^T$.

### D. Adaptation to novel situations

One of the most essential abilities of an LfD method is its adaptability to novel unseen situations.

*1) Adaptation to via-point pose:* Suppose that the robot is asked to pass a via-point pose $g_i^*$ with uncertainty descibed by a covariance matrix $\Sigma_i^*$. The posterior distribution of the trajectory can be computed, which is shown in Fig. 2.

*2) Equivariant adaptation to the change of view:* To change the viewing frame, a group action is applied. Suppose $h \in \text{SE}(3)$ is the relative transformation from the current frame $(O)$ to a new frame $(A)$, then the pose $g$ viewed in frame $O$ can be switched to be viewed in frame $A$ as $g^o = h^{-1}gh$ [9]. The distribution of the whole trajectory as view in frame $A$ can be obtained, which satisfies the equivariance property under the change of view.

*3) Adaptation to robot-specific workspace density:* An important issue for the skill transfer among different robots is the adaptation to the workspace limit and reachability. Previous work has extensively investigated the density of the robot workspaces, in which the more reachable space of the end effector has higher probability [8]. The mathematical foundation

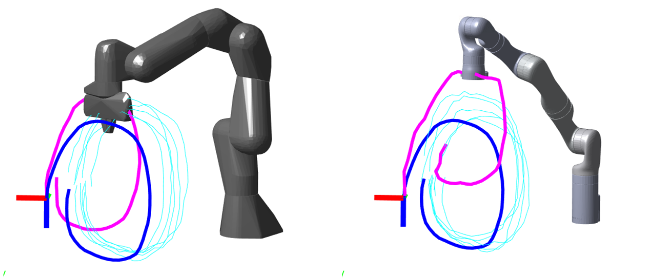

(a) Franka Emika Panda (b) Kinova Gen3

Fig. 3: Fusion with robot-specific workspace density. Thin cyan curves are the demonstrated trajectories using Franka robot; solid thick blue and magenta curves are the conditioned mean trajectory without and with the fusion, respectively.

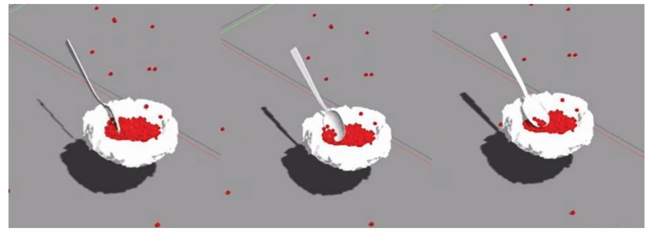

Fig. 4: Robot imagination process for scooping. The container is represented by the white meshed object. A spoon scoops particles that are initially inside the container.

is based on the convolution of Gaussian distributions on $\mathrm{SE}(3)$ [22, 26, 3]. Then, the distribution of each intermediate pose along the trajectory is conditioned by this density function, which can be viewed as a fully observable model. Figure 3 shows the fusion with robot-specific workspace density for Franka Emika Panda and Kinova Gen3 robots.

## III. A Robotic System Combining PRIMP, Workspace-STOMP and Robot Imagination

Common tasks in daily household environment are considered to showcase the proposed robotic system. For each motion primitive, human operators firstly conduct several demonstrations by dragging the robot end effector to fulfill the specific task. The trajectory of the end effector poses for each demonstration is recorded. For a new planning request, *key poses* for the robot are generated from manual inputs, ArUco tags [12] or robot imagination module (as in Sec. III-B). A set of key pose candidates are then fed into PRIMP to condition the trajectory probabilistic distribution. The posterior trajectory distribution tries to reach the desired pose at a certain time step while maintaining the key features of the demonstrations. The learned distribution is then used to guide the STOMP planner with new planning scene, which includes some novel obstacles. Once a feasible trajectory is found by Workspace-STOMP, the robot executes the planned motion the fulfill the designated task. If there is no feasible trajectory, more key pose candidates are generated for re-planning.

### A. Motion planning guided by PRIMP

A novel cost function for the end effector trajectory is proposed to guide the STOMP algorithm, resulting in *Workspace-STOMP*. The cost is computed based on the distance metric in $\mathrm{SE}(3)$ between each rollout trajectory $\mathbf{q}$ at each iteration and the workspace trajectory distribution learned by PRIMP.

The trajectory of the end effector is computed via forward kinematics, denoted as $g(\mathbf{q}, t) \in \mathrm{SE}(3) \times \mathcal{T}$. Then, a number of $m_r$ random samples from the reference trajectory distribution are generated, denoted as $g_r^{(k)} = \left(R_r^{(k)}, \mathbf{t}_r^{(k)}\right)$. And the cost function $\mathbf{c}(\mathbf{q}_i, t_i)$ for $i^{\mathrm{th}}$ time step is computed as

$$\begin{aligned}\mathbf{c}(\mathbf{q}_i, t_i) = & \frac{1}{m_\mathrm{r}} \sum_{k=1}^{m_\mathrm{r}} \left( w_\mathrm{rot} \left\| \log^\vee \left( R^T(\mathbf{q}_i, t_i)\, R_\mathrm{r}^{(k)}(t_i) \right) \right\| \right. , \\ & \left. + w_\mathrm{tran} \left\| \mathbf{t}(\mathbf{q}_i, t_i) - \mathbf{t}_\mathrm{r}^{(k)}(t_i) \right\| \right)\end{aligned} \tag{2}$$

where $w_\mathrm{rot}$ and $w_\mathrm{tran}$ are the weights for rotation and translation parts. The planner is initialized by the mean trajectory of the learned distribution. A plug-in package of the proposed cost function is implemented in MoveIt! platform [10].

### B. Robot imagination

The robot imagination provides functional poses for the pouring and scooping tasks, which are performed using Gazebo physics simulator. The pouring imagination applies the same principle with [27], while the scooping imagination is novel in this work (as shown in Fig. 4). The candidates of the functional poses are sampled uniformly within a range defined by the bounding box of the object of interest. The obtained functional poses are transformed into key poses for the robot and will be used in PRIMP.

## IV. Physical Experiments and Evaluations of the Proposed Methods

### A. Physical experiments

Physical experiments using the Franka Emika Panda robot are conducted. The results for different tasks are demonstrated in Fig. 5.

### B. Benchmarks among learning-from-demonstration methods

The proposed PRIMP method is compared with ProMP [20] and Orientation-KMP [14], both of which are probabilistic LfD methods. For Orientation-KMP, different levels of magnitude for the parameter of kernel are varied. With manually defining 50 different pairs of goal and via-point poses, different LfD algorithms are applied to re-produce the task. The results for the pouring task (Task 1) are shown in Fig. 6.

PRIMP outperforms other probabilistic methods being compared in most cases, in terms of similarity metric for the whole trajectory as well as the distance to the desired via-point pose. This illustrates that PRIMP is able to adapt to a desired via point along the trajectory while maintaining the shape or key features from the demonstration set. The computation is also

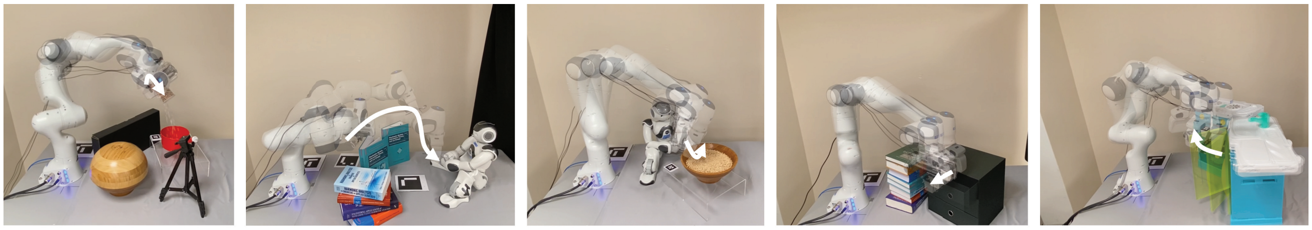

Task 1: Pouring | Task 2: Transporting | Task 3: Scooping | Task 4: Drawer Opening | Task 5: Door Opening

Fig. 5: Physical experiments for different tasks. Planning scenes and key poses are unseen during the demonstration processes.

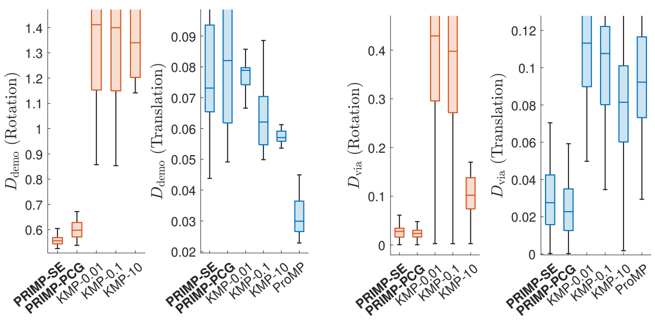


(a) $D_{\text{demo}}$, Task 1 (b) $D_{\text{via}}$, Task 1

Fig. 6: Benchmarks for LfD methods in pouring (Task 1) task. The labels in x-axis are the name of the methods, and the y-axis denotes the comparison metric.

TABLE I: Success rate comparisons among different planners.

| Scene | Task | Cartesian-STOMP | **Workspace-STOMP** |
|---|---|---|---|
| Empty | 2 | 40% | **100%** |
| Sparse | 1 | **10%** | **10%** |
| Cluttered | 4 | 28% | **30%** |
| Narrow | 2 | 12% | **32%** |

efficient, with an averaged time of less than $50\,ms$, which is more than 5 times faster than the compared counterparts.

### C. Benchmarks on guided motion planning

Different planning scenes are constructed using simple geometric primitives in simulation. The proposed Workspace-STOMP planner is compared with the vanilla STOMP [17] and Cartesian-guided STOMP [11]. The initial trajectories for all the planners are set to be the same, referred as a *reference trajectory*. The benchmark results include planning time (Fig. 7), success rate (Tab. I) and distance between the planned and reference trajectories (Fig. 8).

The proposed Workspace-STOMP algorithm runs as fast as the vanilla one, but always has much smaller deviations with the reference trajectory. In more complex and narrow environment, the proposed planner is also competitive with Cartesian-STOMP, which only uses the mean trajectory. The covariance information provided by PRIMP gives more flexibility in varying the samples and guide the optimization through critical regions.

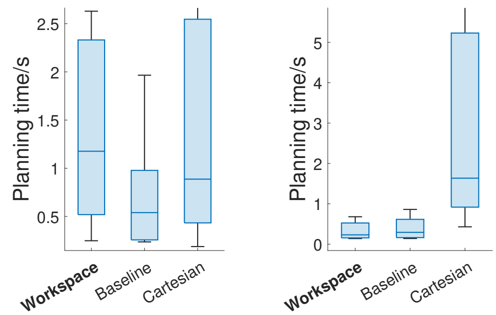


(a) Task 1, narrow scene (b) Task 4, cluttered scene

Fig. 7: Planning time comparisons for different STOMP variants in pouring (Task 1) and drawer opening (Task 4) tasks.

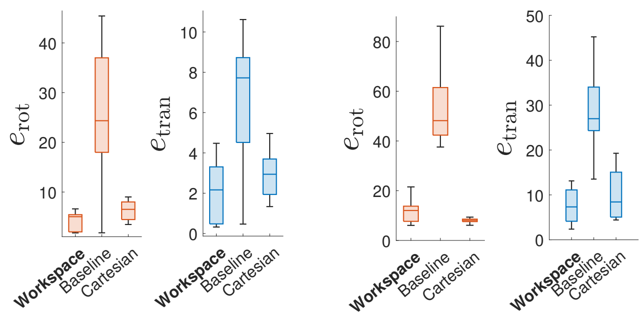


(a) Task 1, narrow scene (b) Task 4, cluttered scene

Fig. 8: Comparisons of distance between planned trajectory and reference trajectory for different variants of STOMP in pouring (Task 1) and drawer opening (Task 4) tasks.

## V. Conclusion

This paper presents *PRobabilistically-Informed Motion Primitives (PRIMP)*, a learning-from-demonstration method that computes the probability distribution in robot workspace. It only requires a few number of or even a single demonstration, and is able to adapt to new via-point poses, a change of viewing frame and robot-specific workspace density. *Workspace-STOMP* planner is then proposed with the guidance of the learned trajectory distribution to avoid novel obstacles. The applicability is demonstrated experimentally in a novel robotic system with the study of object affordance.

## ACKNOWLEDGEMENT

The authors would like to thank Dr. Seng-Beng Ho and Mr. Jikai Ye for useful discussions. This work was supported by NUS Startup grants A-0009059-02-00, A-0009059-03-00, CDE Board account E-465-00-0009-01, National Research Foundation, Singapore, under its Medium Sized Centre Programme - Centre for Advanced Robotics Technology Innovation (CARTIN), sub award A-0009428-08-00 and AME Programmatic Fund Project MARIO A-0008449-01-00.


## REFERENCES


[1] M Kendal Ackerman and Gregory S Chirikjian. A probabilistic solution to the AX=XB problem: Sensor calibration without correspondence. In *International Conference on Geometric Science of Information*, pages 693–701. Springer, 2013.

[2] Sayantan Auddy, Jakob Hollenstein, Matteo Saveriano, Antonio Rodríguez-Sánchez, and Justus Piater. Continual learning from demonstration of robotics skills. *Robotics and Autonomous Systems*, page 104427, 2023.

[3] Timothy D Barfoot and Paul T Furgale. Associating uncertainty with three-dimensional poses for use in estimation problems. *IEEE Transactions on Robotics*, 30 (3):679–693, 2014.

[4] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. Robot programming by demonstration. In *Springer handbook of robotics*, pages 1371–1394. Springer, 2008.

[5] Sylvain Calinon. Gaussians on Riemannian manifolds: Applications for robot learning and adaptive control. *IEEE Robotics & Automation Magazine*, 27(2):33–45, 2020.

[6] Gregory S Chirikjian. Modeling loop entropy. *Methods in enzymology*, 487:99–132, 2011.

[7] Gregory S Chirikjian. *Stochastic models, information theory, and Lie groups, volume 2: Analytic methods and modern applications*, volume 2. Springer Science & Business Media, 2011.

[8] Gregory S Chirikjian and Alexander B. Kyatkin. *Engineering applications of noncommutative harmonic analysis: with emphasis on rotation and motion groups*. CRC press, 2000.

[9] Gregory S Chirikjian, Robert Mahony, Sipu Ruan, and Jochen Trumpf. Pose changes from a different point of view. *Journal of Mechanisms and Robotics*, 10(2), 2018.

[10] Sachin Chitta, Ioan Sucan, and Steve Cousins. Moveit![ROS topics]. *IEEE Robotics & Automation Magazine*, 19(1):18–19, 2012.

[11] Michal Dobiš, Martin Dekan, Adam Sojka, Peter Beňo, and František Duchoň. Cartesian Constrained Stochastic Trajectory Optimization for Motion Planning. *Applied Sciences*, 11(24):11712, 2021.

[12] Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco José Madrid-Cuevas, and Manuel Jesús Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280–2292, 2014.

[13] Seng-Beng Ho. A General Framework for the Representation of Function and Affordance: A Cognitive, Causal, and Grounded Approach, and a Step Toward AGI. *arXiv preprint arXiv:2206.05273*, 2022.

[14] Yanlong Huang, Fares J Abu-Dakka, Joao Silvério, and Darwin G Caldwell. Toward orientation learning and adaptation in Cartesian space. *IEEE Transactions on Robotics*, 37(1):82–98, 2020.

[15] Snehal Jauhri, Jan Peters, and Georgia Chalvatzaki. Robot learning of mobile manipulation with reachability behavior priors. *IEEE Robotics and Automation Letters*, 7(3):8399–8406, 2022.

[16] Wanxin Jin, Todd D Murphey, Dana Kulić, Neta Ezer, and Shaoshuai Mou. Learning from sparse demonstrations. *IEEE Transactions on Robotics*, 2022.

[17] Mrinal Kalakrishnan, Sachin Chitta, Evangelos Theodorou, Peter Pastor, and Stefan Schaal. STOMP: Stochastic trajectory optimization for motion planning. In *2011 IEEE international conference on robotics and automation*, pages 4569–4574. IEEE, 2011.

[18] Dorothea Koert, Guilherme Maeda, Rudolf Lioutikov, Gerhard Neumann, and Jan Peters. Demonstration based trajectory optimization for generalizable robot motions. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 515–522. IEEE, 2016.

[19] Bence Magyar, Nikolaos Tsiogkas, Bruno Brito, Mayank Patel, David Lane, and Sen Wang. Guided stochastic optimization for motion planning. *Frontiers in Robotics and AI*, page 105, 2019.

[20] Alexandros Paraschos, Christian Daniel, Jan Peters, and Gerhard Neumann. Using probabilistic movement primitives in robotics. *Autonomous Robots*, 42(3):529–551, 2018.

[21] Harish Ravichandar, Athanasios S Polydoros, Sonia Chernova, and Aude Billard. Recent advances in robot learning from demonstration. *Annual review of control, robotics, and autonomous systems*, 3:297–330, 2020.

[22] Paul Smith, Tom Drummond, and Kimon Roussopoulos. Computing MAP trajectories by representing, propagating and combining PDFs over groups. In *ICCV*, pages 1275–1282, 2003.

[23] Boyang Ti, Amirreza Razmjoo, Yongsheng Gao, Jie Zhao, and Sylvain Calinon. A geometric optimal control approach for imitation and generalization of manipulation skills. *Robotics and Autonomous Systems*, 164: 104413, 2023.

[24] Nikolaus Vahrenkamp, Tamim Asfour, and Rüdiger Dillmann. Robot placement based on reachability inversion. In *2013 IEEE International Conference on Robotics and Automation*, pages 1970–1975. IEEE, 2013.

[25] Vitalis Vosylius and Edward Johns. Where To Start? Transferring Simple Skills to Complex Environments. In *6th Annual Conference on Robot Learning*, 2022.

[26] Yunfeng Wang and Gregory S Chirikjian. Nonparametric second-order theory of error propagation on motion groups. *The International journal of robotics research*, 27(11-12):1258–1273, 2008.

[27] Hongtao Wu and Gregory S Chirikjian. Can I pour into it? robot imagining open containability affordance of previously unseen objects via physical simulations. *IEEE Robotics and Automation Letters*, 6(1):271–278, 2020.

[28] Gu Ye and Ron Alterovitz. Guided motion planning. In *Robotics research*, pages 291–307. Springer, 2017.

[29] Martijn JA Zeestraten, Ioannis Havoutis, Joao Silvério, Sylvain Calinon, and Darwin G Caldwell. An approach for imitation learning on Riemannian manifolds. *IEEE Robotics and Automation Letters*, 2(3):1240–1247, 2017.

[30] Jihong Zhu, Michael Gienger, and Jens Kober. Learning task-parameterized skills from few demonstrations. *IEEE Robotics and Automation Letters*, 7(2):4063–4070, 2022.