



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Accelerating Motion Planning Via Optimal Transport

An T. Le, Georgia Chalvatzaki, Armin Biess, Jan Peters




“Abundance ~ Discovery”



Planning is reliable 😊



TECHNISCHE
UNIVERSITÄT
DARMSTADT



Urain, J.; Le, A.T.; Lambert, A.; Chalvatzaki, G.; Boots, B.; Peters, J. (2022). Learning Implicit Priors for Motion Optimization, *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
Carvalho, J.; Le, A.T.; Baierl, M.; Koert, D.; Peters, J. (2023). Motion Planning Diffusion: Learning and Planning of Robot Motions with Diffusion Models, *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
Le, A. T.; Hansel, K.; Peters, J.; Chalvatzaki, G. (2023). Hierarchical Policy Blending As Optimal Transport, 5th Annual Learning for Dynamics & Control Conference (L4DC), PMLR.
Le, A. T.; Chalvatzaki, G.; Bliess, A.; Peters, J. (2023). Accelerating Motion Planning via Optimal Transport, *NeurIPS 2023*.

Trajectory Optimization: Collocation method



TECHNISCHE
UNIVERSITÄT
DARMSTADT

$$\boldsymbol{\tau} = [\mathbf{x}_0, \mathbf{u}_0, \dots, \mathbf{x}_{T-1}, \mathbf{u}_{T-1}, \mathbf{x}_T]^\top$$

$$\boldsymbol{\tau}^* = \arg \min_{\boldsymbol{\tau}} \sum_i \lambda_i c_i(\boldsymbol{\tau})$$

$$\text{s.t. } \dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}) \text{ and } \boldsymbol{\tau}(0) = \mathbf{x}_0$$

Model function

(self)-collision avoidance, joint limit, target ee-pose, etc.


Gradient is okay but...




TECHNISCHE
UNIVERSITÄT
DARMSTADT

Trajectory gradients are costly, especially in vectorization settings!


- Need to make sure all costs are differentiable, e.g., obstacle signed distant field
- Dynamics function is also needed to be differentiable




0-level mesh



y plane



z plane




z plane (zoom)



TECHNISCHE
UNIVERSITÄT
DARMSTADT

But how can we solve trajectory optimization without gradients?




Requirements:

- **Batch update**
- **Fast**
- **No gradient access**

$$\min_X f(X) = \min_X \sum_{i=1}^n f(\mathbf{x}_i)$$



TECHNISCHE
UNIVERSITÄT
DARMSTADT




The regular polytopes are unbiased search direction sets!

Sinkhorn Step



Randomly Rotated Polytope Vertices as Step Direction Bases



n optimizing points, m vertices in polytope

$$\mathbf{n} \in \Sigma_n \quad \mathbf{m} \in \Sigma_m$$

$$U(\mathbf{n}, \mathbf{m}) := \{ \mathbf{W} \in \mathbb{R}_+^{n \times m} \mid \mathbf{W} \mathbf{1}_m = \mathbf{n}, \mathbf{W}^\top \mathbf{1}_n = \mathbf{m} \}$$

$n \times m$ cost matrix \mathbf{C}

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \mathbf{S}_k, \quad \mathbf{S}_k = \alpha_k \text{diag}(\mathbf{n})^{-1} \mathbf{W}_\lambda^* \mathbf{D}^P$$

$$\text{s.t. } \mathbf{W}_\lambda^* = \underset{\mathbf{W} \in U(\mathbf{n}, \mathbf{m})}{\text{argmin}} \langle \mathbf{W}, \mathbf{C} \rangle - \lambda H(\mathbf{W})$$

Step vectors (red) are the barycentric projection w.r.t. the polytope family!

Sinkhorn Step



Sinkhorn-Knopp
algorithm

$$\text{OT}_\lambda(\mathbf{n}, \mathbf{m}) := \min_{\mathbf{W} \in U(\mathbf{n}, \mathbf{m})} \langle \mathbf{W}, \mathbf{C} \rangle - \lambda H(\mathbf{W})$$

$$\mathbf{P} = \exp(-\mathbf{C}/\lambda) \quad \mathbf{v}^0 = \mathbf{1}_n$$


Until convergence: $\mathbf{u}^{i+1} = \frac{\mathbf{n}}{\mathbf{P}\mathbf{v}^i}, \quad \mathbf{v}^{i+1} = \frac{\mathbf{m}}{\mathbf{P}^\top \mathbf{u}^{i+1}},$

$$\mathbf{W}_\lambda^* = \text{diag}(\mathbf{u}^*) \mathbf{P} \text{diag}(\mathbf{v}^*)$$


$$\mathbf{X}_{k+1} = \mathbf{X}_k + \mathbf{S}_k, \quad \mathbf{S}_k = \alpha_k \text{diag}(\mathbf{n})^{-1} \mathbf{W}_\lambda^* \mathbf{D}^{\mathbf{P}}$$



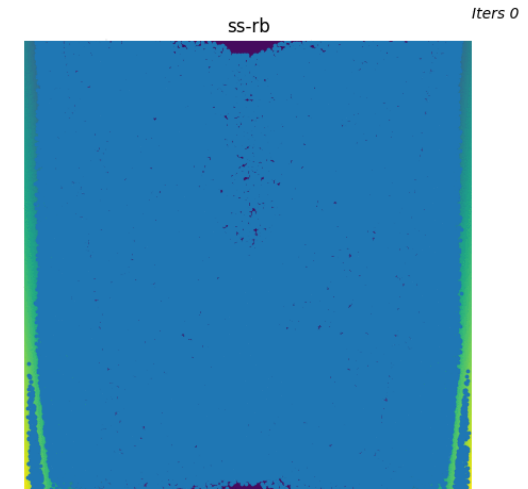
Code available at: <https://github.com/anindex/ssax>




Ackley



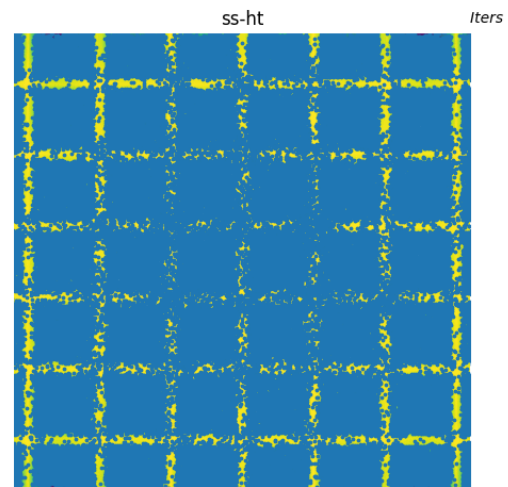
Bukin




Rosenbrock



Egg Holder



Holder Table



Levi




TECHNISCHE
UNIVERSITÄT
DARMSTADT

Now, applying Sinkhorn Step to trajectory optimization?


Anecdote: Go brute-force!



TECHNISCHE
UNIVERSITÄT
DARMSTADT



Iters: 0



We propose Motion Planning via Optimal Transport (MPOT)

- **Massively vectorized planning!**
- **Frame as an optimization problem, then solve it as a sequence of linear programs!**
- **No gradients from task costs or models are required!**

MPOT: Trajectory Optimization




TECHNISCHE
UNIVERSITÄT
DARMSTADT

$$\tau = (X, U) = \{\mathbf{x}_t \in \mathbb{R}^d : \mathbf{x}_t = [\mathbf{x}_t, \dot{\mathbf{x}}_t]\}_{t=0}^T$$

$$\tau^* = \operatorname{argmin}_{\tau} \sum_{t=0}^{T-1} \underbrace{\eta C(\mathbf{x}_t)}_{\text{state cost}} + \underbrace{\frac{1}{2} \|\Phi_{t,t+1} \mathbf{x}_t - \mathbf{x}_{t+1}\|_{\mathbf{Q}_{t,t+1}}^2}_{\text{transition model cost}}$$

**Model constraint
as cost (come from GP prior)**

MPOT: Procedure



1. Construct uniform polytopes with current waypoints as their centers

$$D^P \in \mathbb{R}^{T \times m \times d}$$

2. Populate probing points towards the polytope vertices

$$H^P \in \mathbb{R}^{T \times m \times h \times d}$$

3. Compute local cost matrix

$$C_{t,i} = \frac{1}{h} \sum_{j=1}^h \eta c(\mathbf{x}_t + \mathbf{y}_{t,i,j}) + \frac{1}{2} \|\Phi_{t,t+1} \mathbf{x}_t - (\mathbf{x}_{t+1} + \mathbf{y}_{t+1,i,j})\|_{Q_{t,t+1}^{-1}}^2$$

$$\text{Probe points: } \mathbf{y}_{t,i,j} \in H^P$$

4. Do Sinkhorn Step!

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \mathbf{S}_k, \mathbf{S}_k = \alpha_k \text{diag}(\mathbf{n})^{-1} \mathbf{W}_\lambda^* D^P$$

s.t. $\mathbf{W}_\lambda^* = \underset{\mathbf{W} \in U(\mathbf{n}, \mathbf{m})}{\text{argmin}} \langle \mathbf{W}, \mathbf{C} \rangle - \lambda H(\mathbf{W})$

MPOT



TECHNISCHE
UNIVERSITÄT
DARMSTADT

$$\mathcal{T} = \{\tau_1, \dots, \tau_{N_p}\}$$

$$N = N_p \times T$$

$$D^P \in \mathbb{R}^{N \times m \times d}, H^P \in \mathbb{R}^{N \times m \times h \times d}$$

Algorithm 1: Motion Planning via Optimal Transport

$\mathcal{T}^0 \sim \mathcal{N}(\mu_0, \mathbf{K}_0)$ and $\mathbf{n} = \mathbf{1}_N/N$, $\mathbf{m} = \mathbf{1}_m/m$


while *termination criteria not met* **do**

 (Optional) $\alpha \leftarrow (1 - \epsilon)\alpha$, $\beta \leftarrow (1 - \epsilon)\beta$ // Epsilon Annealing for Sinkhorn Step


 Construct randomly rotated D^P, H^P and compute the cost matrix \mathbf{C} as in Eq. (10)

 Perform Sinkhorn Step $\mathcal{T} \leftarrow \mathcal{T} + \mathbf{S}$

end



MPOT: Experiment



	TF[s]	SUC[%]	S	PL
RRT*	1000 ± 0.00	0	-	-
I-RRT*	1000 ± 0.00	0	-	-
STOMP	-	0	-	-
SGPMP	27.75 ± 0.29	25	0.010 ± 0.001	6.69 ± 0.38
CHOMP	16.74 ± 0.21	40	0.015 ± 0.001	8.60 ± 0.73
GPMP2	40.11 ± 0.08	40	0.012 ± 0.015	8.63 ± 0.53
MPOT	1.49 ± 0.02	55	0.022 ± 0.003	10.53 ± 0.62

This mobile manipulation case has the state space with 36 dimensions!

Key takeaways



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Sinkhorn Step is exciting and needs more theoretical understanding.

- **Solving motion planning with only matrix multiplications 😊**
- **No gradients are required anywhere!**
- **Surprisingly scalability and parallelization capability in massively planning!**
- **Many plans ~ more chance to get better modes!**

Peoples



TECHNISCHE
UNIVERSITÄT
DARMSTADT



An T. Le



Georgia Chalvatzaki



Armin Biess



Jan Peters

Project website:

<https://sites.google.com/view/sinkhorn-step/>

Email: an@robot-learning.de

My website: anthaile.com

I am actively working on Optimal Transport methods applying for Motion Planning and Imitation Learning. Feel free to contact me to hear ranting about Optimal Transport in Robotics 😊