

CS 287, Fall 2015 Optional Extra Credit Problem

Learning from Demonstrations through Trajectory Generalization with Thin-Plate Splines [50pts]

Deliverable: pdf write-up by Friday December 4th, 11:59pm, submitted through Gradescope. **Your pdf should just be a single page. Thanks!**

Please refer to the class webpage for the homework policy. Various starter files are provided on the course website: www.cs.berkeley.edu/~pabbeel/cs287-fa15/.

In this problem we study learning from demonstrations through trajectory generalization with thin-plate splines [1]. Primarily for ease of visualization, we will consider two-dimensional problem settings, but keep in mind this approach equally applies in three-dimensional settings (which would be the typical setting when applying this to robot gripper trajectories).

The approach proceeds as follows: We are given N correspondences, (s_i, t_i) with $s_i \in \mathbb{R}^2$ a point in the source (demonstration) scene, and $t_i \in \mathbb{R}^2$ the corresponding point in the target (test) scene. We seek to find the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that minimizes the following cost:

$$J(f) = \sum_i \|t_i - f(s_i)\|^2 + \lambda \int \|\mathbf{D}^2 f(s)\|_{\text{Frob}}^2 ds. \quad (1)$$

It can be shown that the solution to Eq. (1) can be written as a finite dimensional combination of basis function using the appropriate kernel:

$$f(z) = A^\top K_{\text{tps}}^S(z) + Bz + c \quad (2)$$

Here $A \in \mathbb{R}^{N \times 2}$, $B \in \mathbb{R}^{2 \times 2}$, $c \in \mathbb{R}^{2 \times 1}$, and $K_{\text{tps}}^S(z) \in \mathbb{R}^N$ is the vector of kernel values for each of the points in $S \in \mathbb{R}^{2 \times N}$. In two dimensions, $[K_{\text{tps}}^S(z)]_i = \|s_i - z\|^2 \log \|s_i - z\|$. We will make use of the gram matrix $\mathbf{K}^S \in \mathbb{R}^{N \times N}$, with $\mathbf{K}_{\text{tps}}^S(:, i) = K_{\text{tps}}^S(s_i)$.

The coefficients in A have to satisfy the following constraints:

$$SA = 0, 1^\top A = 0 \quad (3)$$

We can find f by solving the following optimization problem:

$$\begin{aligned} & \underset{A, B, c}{\text{minimize}} && \|T - A^\top \mathbf{K}_{\text{tps}}^S(S) - BS - c1^\top\|_{\text{Frob}}^2 + \lambda \cdot \text{trace}(A^\top \mathbf{K}_{\text{tps}}^S(S)A) \\ & \text{subject to} && SA = 0 \\ & && 1^\top A = 0 \end{aligned} \quad (4)$$

- For A that don't satisfy Eq. 3 we will have $\text{trace}(A^\top \mathbf{K}_{\text{tps}}^S(S)A) < 0$, (i.e. $\mathbf{K}_{\text{tps}}^S(S)$ is not positive semi-definite). Thus, although our optimization problem is convex over the feasible set, the objective is not convex. We need to reformulate this problem before we can use tools like CVX to solve it. Write an equivalent optimization objective that directly includes the constraints in Eq. 3. (*Hint: Is there a change of variables you can use to ensure that the constraints are satisfied?*)
- Implement `compute_warp.m`, which takes S, T , and λ as arguments and returns A, B , and c , the coefficients of the optimal regularized tps warp between the points. There is code to run this and visualize the results in `q1_starter.m`. Three scenarios are provided. The solution to the first scenario is provided for you to check correctness of your implementation. Report the figures showing the warped grids and trajectories for scenarios 2 and 3.

[1] John Schulman, Jonathan Ho, Cameron Lee, Pieter Abbeel, "Learning from Demonstrations through the Use of Non-Rigid Registration." In the proceedings of ISRR 2013.