

Unifying Scene Registration and Trajectory Optimization for Learning from Demonstrations with Application to Manipulation of Deformable Objects

Alex X. Lee Sandy H. Huang Dylan Hadfield-Menell Eric Tzeng Pieter Abbeel

Abstract—Recent work [1], [2] has shown promising results in enabling robotic manipulation of deformable objects through learning from demonstrations. Their method computes a registration from training scene to test scene, and then applies an extrapolation of this registration to the training scene gripper motion to obtain the gripper motion for the test scene. The warping cost of scene-to-scene registrations is used to determine the nearest neighbor from a set of training demonstrations. Then once the gripper motion has been generalized to the test situation, they apply trajectory optimization [3] to plan for the robot motions that will track the predicted gripper motions. In many situations, however, the predicted gripper motions cannot be followed perfectly due to, for example, joint limits or obstacles. In this case the past work finds a path that minimizes deviation from the predicted gripper trajectory as measured by its Euclidean distance for position and angular distance for orientation.

Measuring the error this way during the motion planning phase, however, ignores the underlying structure of the problem—namely the idea that rigid registrations are preferred to generalize from training scene to test scene. Deviating from the gripper trajectory predicted by the extrapolated registration effectively changes the warp induced by the registration in the part of the space where the gripper trajectories are.

The main contribution of this paper is an algorithm that considers this effective final warp as the criterion to optimize for in a unified optimization that simultaneously considers the scene-to-scene warping and the robot trajectory (which were separated into two sequential steps by the past work). This results in an approach that adjusts to infeasibility in a way that adapts directly to the geometry of the scene and minimizes the introduction of additional warping cost. In addition, this paper proposes to learn the motion of the gripper pads, whereas past work considered the motion of a coordinate frame attached to the gripper as a whole. This enables learning more precise grasping motions.

Our experiments, which consider the task of knot tying, show that both unified optimization and explicit consideration of gripper pad motion result in improved performance.

I. INTRODUCTION

Robotic manipulation of deformable objects is challenging due to high-dimensional, continuous state-action spaces and the complicated dynamics of deformable objects. Nevertheless, recent work has shown promising results in enabling robotic manipulation of deformable objects through *learning from demonstrations* [1], [2]. In the first phase, non-rigid registration is used to register a given demonstration scene onto a new test scene, and this registration is extrapolated

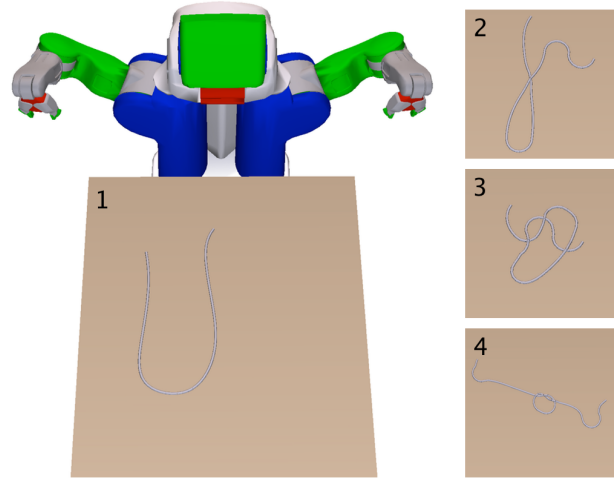


Fig. 1: The overhand knot manipulation task considered in our experiments. A standard knot tie takes three steps, as shown in this particular execution.

to perform *trajectory transfer* of the demonstration’s robot gripper trajectory. Then in the second phase, trajectory optimization is used to generate a feasible sequence of joint angles that produces a gripper trajectory as close as possible to the desired gripper trajectory. The effectiveness of this approach was validated through experiments in knot-tying and suturing.

For complex tasks often demonstrations correspond to steps in the task rather than the entire task itself. Figure 1 shows an example for tying an overhand knot. In general, a single demonstration for a step in the task cannot be expected to cover all possible scenarios that arise during execution. The natural solution to this is to use a library of demonstrations with multiple demonstrations for each step. The trajectory segment from the demonstrations library with the lowest warping cost onto the current scene is chosen.

While experiments with this approach have been promising, thus far they have primarily considered the scenario where the transferred gripper trajectory is feasible and thus executable by the robot. However, the transferred gripper trajectory can often be infeasible, especially in the presence of joint limits or obstacles. When the transferred gripper trajectory is infeasible, past work tries to minimize deviation from this trajectory as measured by Euclidean distance for position and angular distance for orientation.

Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, CA, USA. {alexlee-gk, shhuang, dhm, etzeng, pabbeel}@cs.berkeley.edu

Separately optimizing for a feasible trajectory in this way, however, effectively changes the warp induced by the registration in the part of the space where the gripper trajectories are. This ignores the preference for rigid registrations in generalizing from training scene to test scene. To address this, the main contribution of this paper is a unified optimization that simultaneously considers the scene-to-scene warping and the robot trajectory (which were separated into two sequential steps by the past work). This results in an approach that adjusts to infeasibility in a way that minimizes the introduction of additional warping cost. In addition, this paper proposes to transfer the motion of the gripper pads, whereas past work considered transferring the motion of a coordinate frame attached to the gripper as a whole. Considering the gripper pads explicitly is advantageous because it enables learning more precise grasping motions. Our experiments, which consider the task of knot tying, show that both unified optimization and explicit consideration of gripper pad motion result in improved performance.

II. RELATED WORK

Related work for our contribution stems from learning from demonstrations and from manipulation of deformable objects, in particular knot tying.

The problem of learning from demonstrations deals with the generalization of expert demonstrations to new scenarios [4], [5]. In particular, Calinon et al. [6], [7] developed a method for learning to perform manipulation tasks in different starting scenes. They learn a mixture of Gaussians to represent the environment state and joint angles of the robot across multiple demonstrations, and infer the trajectory for a new environment state by conditioning on that state. Their approach assumes access to a feature representation of the environment, so it is not directly applicable to tasks in environments without fixed feature representations. In contrast, our approach can work directly with point clouds.

In robotics, it is challenging to manipulate deformable objects due to their non-linearity and because the configuration spaces of such objects may be infinite-dimensional [8]. Wada et al. [9] model textile fabric and sponge blocks coarsely and then apply a control method that is robust to discrepancies between the coarse model and the object. Howard et al. [10] present a more general approach for grasping 3D deformable objects that does not assume prior knowledge of the object. They model particle motion of the object using nonlinear partial differential equations, and train a neural network for determining the minimum force required for manipulating the object.

We demonstrate the effectiveness of using unified optimization within trajectory transfer for knot tying, a commonly studied deformable object manipulation task in robotics. Previous approaches to knot tying have demonstrated the utility of rope-specific knowledge and assumptions. For instance, in knot planning from observation (KPO), knot theory is used to recognize rope configurations and define movement primitives from visual observations of humans tying knots [11], [12]. Existing motion planning

approaches for knot tying use topological representations of rope states (i.e., sequences of rope crossings and their properties) and define a model for transitioning between topological states [13], [14], [15]. Robust open loop execution of knot tying has also been explored [16].

Similar to these approaches, our registration will consider transfer only from scenes where the demonstration rope is in the same topological configuration.

III. TECHNICAL BACKGROUND

A. Thin Plate Splines

The thin plate spline (TPS) approach [17] finds a function f mapping between pairs of correspondence points that minimizes the bending cost of f :

$$\begin{aligned} \min_f \quad & \|f\|_{\text{TPS}}^2 \\ \text{subject to} \quad & \mathbf{y}_i = f(\mathbf{x}_i) \end{aligned} \quad (1)$$

where \mathbf{x}_i and \mathbf{y}_i , $i = 1, \dots, N$, refer to pairs of correspondence points in the demonstration and test scenes, respectively. We also refer to this bending cost as warping cost in this paper. For notational convenience, we concatenate the points of the demonstration scene into $X = [\mathbf{x}_1 \dots \mathbf{x}_N]^\top$ and the points of the test scene into $Y = [\mathbf{y}_1 \dots \mathbf{y}_N]^\top$.

The objective $\|f\|_{\text{TPS}}^2$ is the TPS energy function,

$$\|f\|_{\text{TPS}}^2 = \int d\mathbf{x} \|D^2 f(\mathbf{x})\|_F^2, \quad (2)$$

which is a measure of the curvature of the function f .

We can trade off between matching correspondence points and the smoothness of our spline with the following regularized objective

$$f = \arg \min_f \sum_{i=1}^N \|\mathbf{y}_i - f(\mathbf{x}_i)\|_2^2 + \lambda_1 \|f\|_{\text{TPS}}^2, \quad (3)$$

where the parameter λ_1 controls this trade-off.

Duchon showed the minimizer f in Equation (1) can be expressed as a weighted sum of basis functions centered around the data points \mathbf{x}_i , in addition to an affine part [18]. As shown in [17], Equation (3) can be efficiently solved analytically.

B. Learning from Demonstrations with Thin Plate Splines

Schulman et al.'s approach to learning from demonstrations [1] calculates the new trajectory by first using a modified thin plate spline robust point matching (TPS-RPM) algorithm [19] to calculate a non-rigid scene registration that maps points from the demonstration scene to the new scene. TPS-RPM alternates between (1) estimating correspondences between the two scenes' point clouds and (2) fitting the optimal thin plate spline transformation based on these estimated correspondences. Then, the mapping function is applied to the demonstration trajectory in order to obtain a potential trajectory for the new scene. This potential trajectory does not incorporate collision avoidance and joint limits, so trajectory optimization [3] is then used to incorporate these constraints. The hope is that the resulting trajectory will incorporate

variations in the environment and thus succeed in performing the desired manipulation. In practice, this method has been shown to be effective at generalizing expert demonstrations to new, unseen scene configurations.

IV. TPS IMPROVEMENTS AND UNIFIED OPTIMIZATION

In this section, we present improvements on standard TPS trajectory transfer and describe an algorithm for trajectory transfer that directly incorporates feasibility and collision constraints.

A. Improvements on TPS Trajectory Transfer

We propose two changes to the trajectory transfer algorithm from Schulman et al. [1]: an alteration to the TPS optimization that prioritizes matching key points and a change to the way gripper poses are transferred.

In trajectory transfer, not all points need to be transferred with the same accuracy. For points that are close to the demonstration trajectory (e.g., grasping points on an object) we require high accuracy and, in exchange, will accept less rigid mappings in those regions. However, for points far away from the gripper trajectory, high accuracy is not as important and it can make sense to give up accuracy on those points in exchange for performing better on the other terms in the objective. To achieve this, we modify the matching to include explicit constraints on the tightness of fit at points that fall within a minimum distance of the demonstration trajectory.

We also transfer gripper poses from demonstration to test scenes differently from Schulman et al [1]. They use f to transfer the gripper position and use the Jacobian of f to transfer the corresponding gripper orientation for each point in the demonstration trajectory. This can be sensitive to local changes or errors in the mapping function. Instead, we propose to represent gripper poses with a collection of points on the gripper surface. This allows us to transfer the demonstration trajectory by applying f to these representative points, and to then use trajectory optimization to find feasible gripper poses whose representative points best match the warped points. For our experiments, we represent the PR2 gripper pose with a set of points on the finger pads. This change also obviates the need to discretize our gripper into being either open or closed, and potentially applies to more complex gripper demonstration trajectories.

In addition to these two fundamental changes, we use known and fixed correspondences between the two scenes' point clouds to avoid confounding factors when comparing the different trajectory transfer methods. In principle, TPS-RPM could be used to estimate the correspondences, as was done in the previous work [1].

B. Unified TPS and Trajectory Optimization

1) *Motivation:* The motivation behind learning from demonstrations with thin plate splines is that finding a trajectory that minimizes TPS bending cost between the training and test scenes will typically preserve the success of manipulations. In the learning from demonstrations method

from Schulman et al [1] (hereafter referred to as the two-phase method), the result of trajectory transfer may be infeasible, so it is then passed through trajectory optimization. For the manipulation to succeed, the implicit warp of the demonstration trajectory onto the result of trajectory optimization should have low bending cost. However, this is not necessarily the case in the presence of obstacles or underactuated dynamics.

Our goal is to find a *feasible* trajectory that minimizes the bending cost of this implicit warp. The cost we propose includes correspondences between both points in the demonstration and new scenes, and points in the demonstration and new trajectories. The two-phase method can be viewed as performing only one round of block-coordinate descent on this cost in a unified optimization over TPS parameters and robot controls.

2) *Optimization Formulation:* To account for the presence of substantial feasibility constraints, we propose directly minimizing the unified objective. We optimize simultaneously over both the new trajectory, τ , and the transfer function, f :

$$\begin{aligned} \min_{\tau, f} \quad & \alpha C_{\text{TPS}}(X, Y, f) + \beta C_{\text{transfer}}(\tau_d, \tau, f) \\ & + \gamma C_{\text{trajectory}}(\tau) \\ \text{subject to} \quad & \tau \text{ collision-free and feasible} \\ & \|I_Y - f(I_X)\|_2^2 < \epsilon \end{aligned}$$

where τ_d is the demonstrated trajectory, C_{TPS} is the regularized TPS objective, C_{transfer} is a cost on the error between the current trajectory τ and the demonstration trajectory τ_d transferred under f , and $C_{\text{trajectory}}$ is a cost that encodes desired properties of the new trajectory. The variables α , β and γ are constants and they trade off the different costs in the objective. The points I_X , which are a subset of the points X , are points that are required to fall within a given radius of the demonstration trajectory and their correspondences, as discussed in Section IV-A. I_Y consists of the subset of the points Y that correspond to I_X .

C_{TPS} is the thin plate spline cost from Schulman et al.,

$$C_{\text{TPS}}(X, Y, f) = \frac{1}{N} \sum_{i=1}^N \|y_i - f(x_i)\|_2^2 + \lambda_1 \|f\|_{\text{TPS}}^2 + \lambda_2 r(f)$$

where $r(f)$ is a regularizer on the affine part of the transformation f .

The C_{transfer} cost is the error between the representative points of the gripper pads of the current trajectory, $\mathbf{q}_{t,k}$, and the warped representative points of the demonstration trajectory, $\mathbf{p}_{t,k}$, for every time step t and for every representative point k on the gripper:

$$C_{\text{transfer}}(\tau_d, \tau, f) = \frac{1}{T} \sum_{t=1}^T \sum_{k=1}^K \|\mathbf{q}_{t,k} - f(\mathbf{p}_{t,k})\|_2^2$$

where $\mathbf{q}_{t,k}$ and $\mathbf{p}_{t,k}$ are deterministic functions of τ and τ_d , respectively.

$C_{\text{trajectory}}$ encodes desired properties of the trajectory. In our case, we penalize for changes of each joint angle j to encourage shorter paths:

$$C_{\text{trajectory}}(\tau) = \frac{1}{T-1} \sum_{t=1}^{T-1} \sum_{j=1}^J (\theta_{t+1,j} - \theta_{t,j})^2.$$

We minimize this objective with the penalty SQP method described in [3]. This is slower than standard trajectory optimization, but a customized algorithm could improve on these results. We will explore this in future work.

V. EXPERIMENTS AND RESULTS

Our quantitative experiments evaluate the impact of using improved TPS trajectory transfer and unified optimization on the performance of tying an overhand knot when the robot’s degrees of freedom (DOF) are reduced. We use the Willow Garage PR2 robot in a simulation environment that uses Bullet Physics [20] as the physics engine, with the rope represented as a linked chain of capsules with constraints on bending and torsion.

We also illustrate the intuition behind our approach on a simple example with a holonomic car, in which the original warped trajectory is infeasible due to a new obstacle in the test scene. Unified optimization finds a qualitatively better trajectory for the test scene than the two-phase method.

A. Knot Tying with the PR2

1) *Demonstration Library*: We use a set of 148 demonstrations collected by Schulman et al. [1] to generalize from. Each demonstration step consists of a gripper trajectory and a point cloud of the state of the rope at the start of the step. Demonstration gripper trajectories were collected kinesthetically, capturing the pose of the grippers at each time step as a human moves them in tying a knot. The point clouds were collected by an RGBD camera, filtering out all points that are not part of the rope (e.g., points of the table). This dataset consists of demonstrations corresponding to steps of a three-step method of tying an overhand knot (Figure 1) and demonstrations of recovery from common failure cases.

2) *Benchmark*: Our benchmarks are drawn from a distribution of initial rope states, generated by perturbing the initial rope states of the demonstration library. Perturbing a rope consists of selecting five uniformly spaced points along an original initial rope state, and dragging each in a random direction within a specified radius. In our experiments, we use a benchmark with a perturbation radius of 10 cm and a harder benchmark with a radius of 15 cm.

Each benchmark consists of 100 initial rope states. A task is considered successful if the robot ties an overhand knot within five or fewer steps, starting from the given initial test state. Aside from the rope configuration, the test environment is identical to the demonstration environment: the rope is on a table with no obstacles. However, in the test scenarios we impose stricter angle limits for the arms’ joints, to increase the likelihood of scenarios where the original warped trajectory is infeasible. We reduce the range of each

arm’s joint angle limits by several factors; the reduced range is centered around the midpoint between the original lower and upper limits.

At each step in tying an overhand knot, the demonstration with the lowest warping cost is selected to generalize from. Sometimes the trajectory optimization results in trajectories where some robot parts collide with the robot itself or the table, even though trajectory optimization has a cost for avoiding collisions with a safety distance of 2.5 cm. This can occur because the trajectory optimization reached an undesirable local optima, or it may occur because the collision cost was outweighed by the other costs in the objective. When this happens, we try generalizing from the demonstration with the next lowest warping cost, and repeat until either the trajectory is feasible or a maximum of 10 demonstrations have been tried. The trajectory is feasible if the joint angle trajectory upsampled to 100 time steps is collision free at each of those time steps. We consider all these trials as part of the same step, since collision checking can be done without executing the trajectory.

3) *Experiments*: On these two benchmarks, we evaluate (1) the improved TPS trajectory transfer method and (2) the unified optimization method. As a baseline, we compare our results against the original trajectory transfer method introduced by Schulman et al. [1], with some minor differences.

Their method uses a modified robust point matching algorithm, TPS-RPM, to estimate correspondences and fits a bidirectional thin plate spline between the demonstration and test scene’s rope point clouds, regardless of the underlying topological rope configuration.

In contrast, we use known and fixed correspondences, fit a unidirectional thin plate spline, and only consider registrations from demonstration scenes where the starting configuration of the rope is topologically equivalent to the rope configuration of the current test scene. For this, our point clouds have an underlying structure: the rope points are ordered and evenly spaced points along the rope’s backbone. These points were manually labeled from the RGBD images of demonstration scenes, and they were available from the simulation in the test scenes. The fixed correspondences were determined by matching topological rope segments between the ropes of both scenes and resampling the number of points in each segment of the test rope to match that of the corresponding segment of the demonstration rope. More concretely, given a point cloud \hat{Y} with \hat{N} points in the test scene, the correspondence and resampling can be represented by a $N \times \hat{N}$ matrix M , such that $Y = M\hat{Y}$ and there is a one-to-one correspondence between points \mathbf{x}_i and \mathbf{y}_i .

In our experiments, I_X are the set of points in X that are within 5 cm of the gripper’s fingers when they grasp the rope in the demonstration. We used an ϵ of 5 mm, which is also the radius of the simulated rope.

4) *Results*: The success rates of these three trajectory transfer methods are shown in Tables Ia and Ib for the benchmarks with 10 cm and 15 cm perturbations, respectively.

As expected, we see that in general the improved TPS trajectory transfer and unified optimization perform better

than the previous method. In addition, unified optimization does significantly better than the other two methods in cases where the original warping function generalizes to target gripper poses or finger pad points that are infeasible due to joint angle constraints. This is especially apparent when the range of the degrees of freedom has been reduced by a factor of 0.4. Beyond a factor of 0.4, the trajectory optimization of all the three methods becomes significantly harder to solve, and most tasks can no longer be executed successfully.

The time taken to select among the 148 demonstrations per step across all the experiments was 0.89 s on average with a standard deviation of 0.41 s. The execution times for optimizing the transferred trajectory and simulating it per step across all the experiments are shown in Table II. The computation time performance was evaluated using a combination of Python and C++ implementations running on a single 3.5 GHz Intel i7 processor core. The reported times are *real* times. Although single experiments were run on single cores, multiple experiments were run simultaneously on a single machine, so the reported *real* times are just an upper bound on the computation *CPU* times.

B. Holonomic Car

To provide intuition for why our unified optimization leads to improved performance, we consider a toy example involving a 2D holonomic car. The demonstration trajectory is a driving path from left to right in a straight line. An obstacle is introduced in the middle of the new scene, blocking the trajectory. Figure 2a shows the ground-truth warp from the demonstration to the new scene, given pre-defined correspondence points between the two scenes. The trajectory is warped upward, but is blocked by the obstacle.

The two-phase approach (shown in Figure 2b) poorly navigates around the obstacle, following its boundary closely. On the other hand, the unified optimization approach (shown in Figure 2c) smoothly navigates around the obstacle. This occurs because unified optimization can influence the warp to direct the trajectory *above* the obstacle, rather than colliding with it. This qualitatively demonstrates that unified optimization produces better-conditioned trajectories.

VI. CONCLUSION

We considered the problem of generalizing demonstrated motions to new, yet similar, situations. We built on prior work that consists of a two-phase approach: in the first phase a registration between training and test scene is computed, and this registration is extrapolated to transfer the gripper motion from training situation to test situation. In the second phase optimization-based motion planning finds a motion for the robot to follow the transferred gripper motion.

In this paper we described an algorithm for performing a unified optimization that simultaneously considers the scene-to-scene warping (from the first phase) and the robot trajectory (from the second phase). This results in an approach that adjusts to trajectory infeasibility in a way that adapts directly to the geometry of the scene and minimizes the introduction of additional warping cost. In addition, this paper proposed

DOF Range Reduction Factor	Schulman et al. [1]	Improved TPS Trajectory Transfer (Section IV-A)	Unified TPS and Trajectory Optimization (Section IV-B)
1.0	87%	99%	99%
0.9	84%	96%	98%
0.8	82%	96%	99%
0.7	83%	96%	100%
0.6	87%	98%	100%
0.5	88%	88%	98%
0.4	62%	51%	83%

(a) Success rates for the benchmark with 10 cm initial state perturbations.

DOF Range Reduction Factor	Schulman et al. [1]	Improved TPS Trajectory Transfer (Section IV-A)	Unified TPS and Trajectory Optimization (Section IV-B)
1.0	88%	91%	92%
0.9	85%	93%	94%
0.8	87%	93%	92%
0.7	91%	93%	93%
0.6	85%	93%	92%
0.5	80%	83%	93%
0.4	55%	47%	71%

(b) Success rates for the benchmark with 15 cm initial state perturbations.

TABLE I: Success rates of tying a knot using different approaches for transferring the demonstration trajectory. The approach of Schulman et al. registers the demonstration scene onto the new scene using a TPS and then uses trajectory optimization to minimize the error between a feasible sequence of gripper poses and the warped demonstration gripper poses. The improved TPS trajectory transfer approach registers both scenes using a TPS that constrains the scene points near the grasping point and then uses trajectory optimization to minimize the error between the points of the gripper pads and the warped points of the demonstration. The unified optimization approach penalizes for the same TPS and trajectory objectives as in the last method, but jointly optimizes for both the TPS and the trajectory in the same optimization problem.

	Schulman et al. [1]	Improved TPS Trajectory Transfer (Section IV-A)	Unified TPS and Trajectory Optimization (Section IV-B)
Mean Time	3.81 s	6.89 s	25.2 s
Standard Deviation Time	3.07 s	6.52 s	33.0 s

TABLE II: Execution times for optimizing the transferred trajectory and simulating it per step across all the experiments.

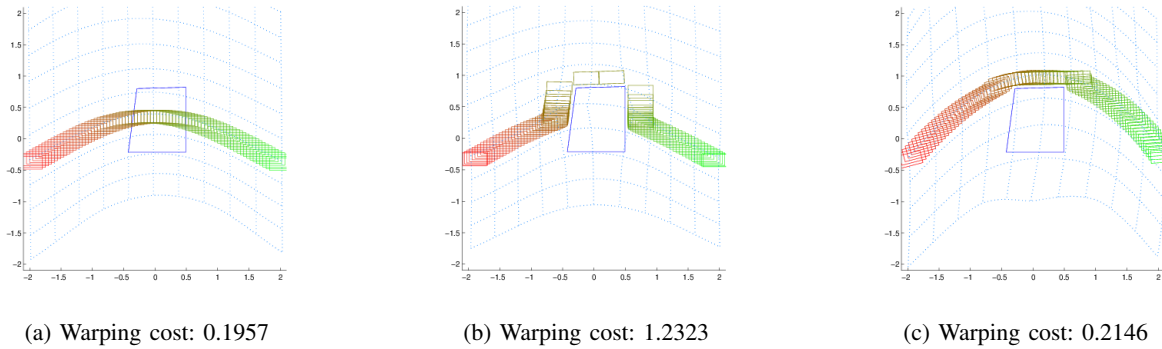


Fig. 2: We qualitatively compared the two-phase optimization method to our unified optimization method in a holonomic car example, in which a new obstacle is introduced in the test scene. The demonstration trajectory is a driving path from left to right in a straight line. The warping function is visualized as a warped grid. (a) The pre-defined correspondence points between the demonstration and test scenes finds a warping function that bends upwards; the warped trajectory is in collision with the obstacle. (b) After enforcing feasibility constraints, the original two-phase method produces a trajectory that poorly navigates the obstacle. (c) Unified optimization produces a trajectory that smoothly avoids the obstacle. The warping costs for (b) and (c) refer to that of the resulting trajectory after optimization when using the car corner points as correspondences.

to learn the motion of the gripper pads, whereas past work considered the motion of a coordinate frame attached to the gripper as a whole. This new proposed approach enables learning more precise grasping motions.

Our experiments, which consider the task of knot tying, show that both unified optimization and explicit consideration of the gripper pad motion result in improved performance. Unified optimization performs significantly better in cases where the warped demonstration gripper trajectories are further away from trajectories achievable by the robot.

VII. ACKNOWLEDGMENTS

We thank George Mulcaire for insightful discussions and advice. This research has been funded in part by DARPA Young Faculty Award #D13AP00046 and by a Sloan Fellowship. Alex Lee has been funded by an NSF Graduate Research Fellowship, Sandy Huang by a Berkeley Fellowship, and Dylan Hadfield-Menell by a Chancellor's Fellowship.

REFERENCES

- [1] J. Schulman, J. Ho, C. Lee, and P. Abbeel, "Learning from Demonstrations through the Use of Non-Rigid Registration," in *Proceedings of the 16th International Symposium on Robotics Research (ISRR)*, 2013.
- [2] J. Schulman, A. Gupta, S. Venkatesan, M. Tayson-Frederick, and P. Abbeel, "A Case Study of Trajectory Transfer Through Non-Rigid Registration for a Simplified Suturing Scenario," in *Proceedings of the 26th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013.
- [3] J. D. Schulman, J. Ho, A. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Proceedings of Robotics: Science and Systems (RSS)*, 2013.
- [4] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A Survey of Robot Learning from Demonstration," *Robot. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, May 2009.
- [5] S. Schaal, "Is imitation learning the route to humanoid robots?" *Trends in cognitive sciences*, vol. 3, no. 6, pp. 233–242, 1999.
- [6] S. Calinon, F. Guenter, and A. Billard, "On Learning, Representing, and Generalizing a Task in a Humanoid Robot," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 37, no. 2, pp. 286–298, April 2007.
- [7] S. Calinon, F. D'halluin, D. Caldwell, and A. Billard, "Handling of multiple constraints and motion alternatives in a robot programming by demonstration framework," in *Humanoid Robots, 2009. Humanoids 2009. 9th IEEE-RAS International Conference on*, Dec 2009, pp. 582–588.
- [8] F. Lamiraux and L. E. Kavraki, "Planning Paths for Elastic Objects Under Manipulation Constraints," *International Journal of Robotics Research*, vol. 20, pp. 188–208, 2001.
- [9] T. Wada, S. Hirai, H. Mori, and S. Kawamura, "Robust Manipulation of Deformable Objects Using Model Based Technique," in *Articulated Motion and Deformable Objects*, ser. Lecture Notes in Computer Science, H.-H. Nagel and F. Perales Lpez, Eds. Springer Berlin Heidelberg, 2000, vol. 1899, pp. 1–14.
- [10] A. M. Howard and G. A. Bekey, "Intelligent Learning for Deformable Object Manipulation," *Autonomous Robots*, vol. 10, no. 1, pp. 51–58, 2000.
- [11] T. Morita, J. Takamatsu, K. Ogawara, H. Kimura, and K. Ikeuchi, "Knot planning from observation," in *Proceedings of the International Conference on Robotics and Automation (ICRA)*. IEEE, 2003, pp. 3887–3892.
- [12] J. Takamatsu, T. Morita, K. Ogawara, H. Kimura, and K. Ikeuchi, "Representation for Knot-tying Tasks," *Trans. Rob.*, vol. 22, no. 1, pp. 65–78, Nov. 2006. [Online]. Available: <http://dx.doi.org/10.1109/TRO.2005.855988>
- [13] M. Moll and L. E. Kavraki, "Path planning for deformable linear objects," *IEEE Transactions on Robotics*, vol. 22, pp. 625–636, 2006.
- [14] M. Saha, P. Ito, and J.-C. Latombe, "Motion Planning for Robotic Manipulation of Deformable Linear Objects," in *Experimental Robotics*, ser. Springer Tracts in Advanced Robotics, O. Khatib, V. Kumar, and D. Rus, Eds. Springer Berlin Heidelberg, 2008, vol. 39, pp. 23–32.
- [15] H. Wakamatsu, E. Arai, and S. Hirai, "Knotting/Unknotting Manipulation of Deformable Linear Objects," *The International Journal of Robotics Research*, vol. 25, no. 4, pp. 371–395, 2006.
- [16] M. Bell, "Flexible Object Manipulation," Ph.D. dissertation, Dartmouth College, Hanover, NH, USA, 2010.
- [17] G. Wahba, *Spline Models for Observational Data*. Philadelphia: Society for Industrial and Applied Mathematics, 1990.
- [18] J. Duchon, "Splines minimizing rotation-invariant semi-norms in sobolev spaces," in *Constructive Theory of Functions of Several Variables*, ser. Lecture Notes in Mathematics, W. Schempp and K. Zeller, Eds. Springer Berlin Heidelberg, 1977, vol. 571, pp. 85–100. [Online]. Available: <http://dx.doi.org/10.1007/BFb0086566>
- [19] H. Chui and A. Rangarajan, "A new point matching algorithm for non-rigid registration," *Computer Vision and Image Understanding*, vol. 89, no. 2-3, pp. 114–141, 2003.
- [20] E. Coumans, "Bullet Physics," Jan 2014. [Online]. Available: <http://code.google.com/p/bullet/>