Using Gaussian Process Regression for Efficient Motion Planning in Environments with Deformable Objects

Barbara Frank and **Cyrill Stachniss** and **Nichola Abdo** and **Wolfram Burgard**

Institute of Computer Science
Autonomous Intelligent Systems
Georges-Koehler-Allee 079
D-79110 Freiburg, Germany

Abstract

The ability to plan their own motions and to reliably execute them is an important precondition for autonomous robots. In this paper, we consider the problem of planning the motion of a mobile manipulation robot in the presence of deformable objects in the environment. Our approach combines probabilistic roadmap planning with a deformation simulation system. Since the physical deformation simulation is computationally demanding, we use an efficient variant of Gaussian process regression to estimate the deformation cost for individual objects based on training examples. We generate the training data by employing a simulation system in a preprocessing step. Consequently, no simulations are needed during runtime. We implemented and tested our approach on a mobile manipulation robot. Our experiments show that the robot is able to accurately predict and thus consider the deformation cost its manipulator introduces to the environment during motion planning. Simultaneously, the computation time is substantially reduced compared to a system that performs physical simulations online.

1 Introduction

The ability to plan its own motion is an important capability of a truly autonomous robot. There is a large body of literature on path and motion planning for mobile robots, most of them assume a static world or environments that consist of rigid objects only. Recently, several researchers addressed the problem of planning for deformable robots (Holleman, Kavraki, and Warren 1998; Anshelevich et al. 2000; Bayazit, Lien, and Amato 2002; Gayle et al. 2005) or the problem of dealing with deformable environments (Rodríguez, Lien, and Amato 2006) or deformable objects such as cloth or towels (Maitin-Shepard, M. Cusumano-Towner, and Abbeel 2010). Real world applications of planning in deformable environments include surgical simulations, where the interaction with (and potential injury of) organs should be minimized (Gayle et al. 2005; Maris, Botturi, and Fiorini 2010).

A straightforward way of considering deformations of objects during planning is to generate collision-free trajectories while considering all deformable objects as free space. During path planning, the planner has to simulate the deformation of the objects resulting from the interaction with
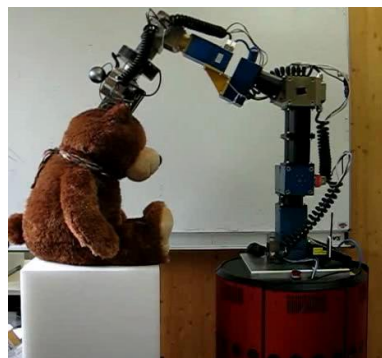

Figure 1: Our mobile manipulation robot Zora deforming a plush teddy bear.

the robot and its manipulator and consider these additional costs. The problem with this method is that an appropriate physical simulation typically requires substantial computational resources, which makes such an approach unsuitable for realistic problems. In our previous work (Frank et al. 2008; 2009), we considered the problem of 2D navigation among deformable objects in the context of reactive collision avoidance systems. Our method approximated the deformation cost function using a low-dimensional grid to allow for an efficient estimation of the expected deformation cost. This approach, however, is no longer feasible for manipulation robots due to their typically high-dimensional configuration spaces.

In this paper, we present a novel approach that applies efficient Gaussian process regression to approximate the deformation cost functions of objects in the configuration space of the robot. This allows to efficiently plan trajectories in the presence of deformable objects even for manipulation robots such as the one shown in Fig. 1. Throughout this paper we assume that the robot can deform the objects but cannot move them in the environment. To improve the efficiency of the learning process, we furthermore sample a restricted set of trajectories only. In different experiments, we demonstrate that our approach yields accurate estimates and, at the same time, allows for efficient planning of trajectories along which the robot interacts with deformable objects.

2 Related Work

Recently, several path planning approaches for deformable robots in static environments have been presented (Holleman, Kavraki, and Warren 1998; Anshelevich et al. 2000; Bayazit, Lien, and Amato 2002; Gayle et al. 2005). These approaches have in common that a probabilistic roadmap is applied to plan motions and a deformation simulation is used to compute the expected deformations. The considered deformation models employed in the different approaches vary. Robots are assumed to be surface patches (Holleman, Kavraki, and Warren 1998) or consist of primitive volumetric elements (Anshelevich et al. 2000) and are modeled using mass-spring systems. Gayle et al. (2005) add constraints for volume preservation to achieve a physically more realistic simulation of deformations. In contrast to our approach, these planners deform the robot rather than the obstacles to avoid collisions. Rodríguez et al. (2006) proposed an approach to planning in completely deformable environments. They employ a mass-spring system with additional physical constraints for volume-preservation to enforce a more realistic behavior of deformable objects. They use rapidly exploring random trees and apply virtual forces to expand the leaves of the tree until the goal state is reached. An approach presented by Maris et al. (2010) plans paths for a surgical tool. In this work, the organs are modeled as deformable objects and the aim is to minimize their deformation as well as penetration. This is done by optimizing the control points of a path with respect to constraints that consider the stiffness of objects and the penetration depth of the surgical tool. The tool, however, is constrained to a rod, that always has to pass through a fixed point (the insertion position), and the degrees of freedom are limited to four.

A drawback of the approaches discussed above is that they need to compute the deformation simulations during runtime. To deal with the high computational load for real robots, we presented an approximation of the deformation cost function for wheeled robots moving in a plane (2008; 2009) that can run online. In our new work, we extend our previous approach to the more complex problem of planning motions for manipulators that operate in 3D. In this setting, the possible trajectories that need to be considered are more complex and thus a more sophisticated method for estimating the deformation costs is needed. We present an efficient approximation based on Gaussian processes that allows to carry out motion planning tasks on the fly.

In the context of robot learning tasks, Gaussian processes (GPs) are becoming increasingly popular. A good introduction into GPs can be found in (Rasmussen and Williams 2006). In robotics, GPs have been used e. g. for terrain modeling (Vasudevan et al. 2009), learning motion and observation models (Ko and Fox. 2009) and several other problems. In some parts, the approach of Vasudevan et al. (2009) is similar to our method. To model large outdoor terrain structures, they perform a nearest-neighbor query on measured elevation data and consider only inputs in the local neighborhood of the query point. This is done efficiently using a KD-tree. We apply the same trick to reduce the number of training samples for the GP to the subset of the most relevant ones for solving the regression problem at hand.

3 Motion Planning in the Presence of Deformable Objects

3.1 Planning using Probabilistic Roadmaps

To plan trajectories for our manipulation robot, we use the probabilistic roadmap framework (Kavraki et al. 1996). The key idea is to represent the collision-free configuration space of the robot by a set of samples that form the nodes of a graph. Edges in this graph model feasible trajectories between neighboring configurations. Such a roadmap can be precomputed given a model of the environment. To actually plan a trajectory for the robot, one connects the current robot configuration as well as the target configuration to the graph. Most motion planning systems assign costs to the edges that correspond to their distance in configuration or work space. Then, this graph allows for applying graph search techniques such as $A^\star$ or Dijkstra's algorithm to search for the optimal path between a given starting and goal point in the roadmap.

Since we are interested in considering deformable objects, we also allow for samples and edges that lead to collisions with these objects when generating the probabilistic roadmap. Accordingly, we need to consider the deformation costs when planning trajectories. Our system uses a weighted sum between the distance of the nodes in configuration space and the deformation costs. For an edge between the nodes i and j, its cost is given by

$$C(i,j) \quad := \quad \alpha \, C_{def}(i,j) + (1-\alpha) \, dist(i,j), \quad (1)$$

where $\alpha \in [0,1]$ is a user-defined weighting coefficient. The term $C_{def}(i,j)$ represents the costs that are introduced when the robot deforms objects along its trajectory and the term $dist(i,j)$ corresponds to the distance between nodes in configuration space. Our current implementation applies $A^\star$ to find the optimal path in the roadmap given Eq. (1). To obtain an admissible heuristic for $A^\star$, i. e., a heuristic that underestimates the real costs, we use the distance to the goal configuration weighted with $(1-\alpha)$. Thus, we are able to find the path in the roadmap that optimizes the trade-off between travel cost and deformation cost.

The key difficulty when considering deformable objects in real world planning tasks is to obtain the cost of deformations, i. e., estimating the term $C_{def}(i,j)$, in an efficient way. One possible way to determine this quantity is to perform a physical simulation of the robot movement.

3.2 Determining Deformation Costs

To determine the object deformations introduced by the robot and the associated costs, we employ a physical simulation engine that is based on finite element methods (FEM). In particular, we use DefCol Studio (Heidelberger et al. 2006) as our simulation environment. It combines an FEM-based simulation of the deformations on volumetric meshes following the approaches described in (Hauth and Strasser 2004; Mueller and Gross 2004), with an efficient collision handling scheme. In our previous work, we presented an approach for building such meshes from sensor data and estimating the deformation parameters for real objects (Frank et al. 2010). The parameters, which cannot be observed directly, are estimated by actively deforming a real object
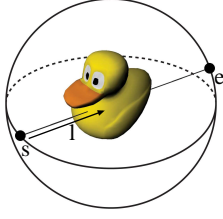
Figure 2: Trajectory parametrization: starting point s and end point e on a virtual sphere around the deformable object together with the distance l from s towards the object.

and by optimizing the deformation parameters in simulation until the observed and the simulated deformation match. Here, we use the parameters estimated with our previous method (Frank et al. 2010).

To define a measure for the deformation cost introduced by the robot, we use the inner potential energy of the object, which is computed in each simulation timestep based on the external forces (resulting from collisions with the robot) and corresponds to the deformation of the object. The deformation cost of a trajectory hence is the sum of the deformation costs of all objects that are in collision with the robot over all timesteps. For further details, we refer the reader to our previous work (Frank et al. 2008).

3.3 Limitations

The approach described so far can be used for planning the trajectory of a robot and its manipulator among deformable objects. The key problem, however, is the computational requirements. Although the deformation simulation can be executed online for a scene, a large number of trajectory hypotheses needs to be evaluated for building the roadmap as well as for planning a trajectory using $A^{\star}$. Additionally, small changes in the world require to recompute the costs for the edges of the roadmap—this makes real world applications basically impossible. To overcome this limitation, the next section presents an efficient way to accurately estimate the deformation costs for individual objects using Gaussian process regression. Our approach uses the simulation system to generate the training inputs and estimates the deformation costs for new trajectories or in a modified environment based on the training data that are generated beforehand. The combination of the planning system and the regression technique allows for efficient planning among deformable objects.

4 Efficient Estimation of the Deformation Cost using Gaussian Process Regression

4.1 Parametrization

The problem of estimating the deformation cost introduced by a robot can be efficiently approached by regression techniques. Let $y_{1:n}$ be the deformation cost values obtained from simulation where the virtual robot executed n different trajectories $\mathbf{x}_{1:n}$. Then, the goal is to learn a predictive model $p(y_* \mid \mathbf{x}_*, \mathbf{x}_{1:n}, y_{1:n})$ for estimating the deformation cost y_* given a query trajectory $\mathbf{x}_*$.

In theory, all possible trajectories through a deformable object can be executed. To bound the complexity of the regression problem, we consider only straight line motions through the object. This is an assumption but not a really strong one since the trajectories generated by most roadmap planners are often piecewise linear motions. The motions considered to estimate the deformation cost are parametrized by five parameters: a starting point s and end point e on a virtual sphere around the object. The points s and e are each described by an azimuth ϕ and an elevation angle θ, together with a distance l from the starting point that describes the length of the motion. Fig. 2 illustrates this parametrization. Thus, $\mathbf{x}_i$ is a five-dimensional vector in our case with $\mathbf{x}_i = [\theta_i^s, \phi_i^s, \theta_i^e, \phi_i^e, l_i]^T$.

4.2 Regression for Estimating Deformation Costs

We approach the problem of estimating the deformation costs by means of nonparametric regression using the Gaussian process (GP) model (Rasmussen and Williams 2006). In this Bayesian approach to non-linear regression, one places a prior on the space of functions using the following definition: A Gaussian process is a collection of random variables, any of which have a joint Gaussian distribution. More formally, if we assume that $\{(\mathbf{x}_i, f_i)\}_{i=1}^n$ with $f_i = f(\mathbf{x}_i)$ are samples from a Gaussian process and define $\mathbf{f} = (f_1, \ldots, f_n)^\top$, we have

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K}), \quad \boldsymbol{\mu} \in \mathbb{R}^n, \mathbf{K} \in \mathbb{R}^{n \times n} . \tag{2}$$

For simplicity, we set $\boldsymbol{\mu} = \mathbf{0}^1$. The interesting part of the GP model is the covariance matrix $\mathbf{K}$. It is specified by $[\mathbf{K}]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ using a covariance function k. Intuitively, the covariance function specifies how similar two function values $f(\mathbf{x}_i)$ and $f(\mathbf{x}_j)$ are. The standard choice for k is the squared exponential covariance function

$$k_{SE}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp\left(-\frac{1}{2} \frac{|\mathbf{x}_i - \mathbf{x}_j|^2}{\ell^2}\right), \tag{3}$$

where the so-called length-scale parameter ℓ defines the global smoothness of the function f and σ_f^2 denotes the amplitude (or signal variance) parameter. These parameters, along with the global noise variance σ_n^2 that is assumed for the noise component, are known as the hyperparameters of the process.

The standard squared exponential covariance function given in Eq. (3) is clearly suboptimal for our problem. The reason for that is our parametrization, which is based on four angles and one Euclidean distance. Considering these dimensions alike does not allow us to model the "similarity" between trajectories well. Therefore, we define a variant of the squared exponential covariance function that considers that these angles are used to describe two points on a sphere. Thus, we consider the distance between the starting points and the end points lying on the sphere from the two inputs

---

[1]The expectation is a linear operator and for any deterministic mean function $m(\mathbf{x})$, the Gaussian process over $f'(\mathbf{x}) := f(\mathbf{x}) - m(\mathbf{x})$ has zero mean.

$\mathbf{x}_i$ and $\mathbf{x}_j$ plus the difference in the length of the trajectory. This results in

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp\left(-\frac{1}{2}\frac{d^2(\mathbf{x}_i, \mathbf{x}_j)}{\ell^2}\right), \qquad (4)$$

with

$$
\begin{aligned}
d(\mathbf{x}_i, \mathbf{x}_j) &= \|l_i - l_j\| + \|\mathrm{p2e}(\theta_i^s, \phi_i^s) - \mathrm{p2e}(\theta_j^s, \phi_j^s)\| + \\
&\quad \|\mathrm{p2e}(\theta_i^e, \phi_i^e) - \mathrm{p2e}(\theta_j^e, \phi_j^e)\|
\end{aligned} \qquad (5)
$$

and where $\mathrm{p2e}(\cdot)$ is the mapping of the spherical coordinates to points on the sphere in $\mathbb{R}^3$.

Given a set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ of training data obtained from the simulation engine, we aim at predicting the target value y_* for a new trajectory specified by $\mathbf{x}_*$. Let $\mathbf{X} = [\mathbf{x}_1; \ldots; \mathbf{x}_n]^\top$ be the matrix of the inputs and $\mathbf{X}_*$ be defined analogously for multiple test data points. In the GP model, any finite set of samples is jointly Gaussian distributed. To make predictions at $\mathbf{X}_*$, we obtain the predictive mean

$$\mathbb{E}[f(\mathbf{X}_*)] = k(\mathbf{X}_*, \mathbf{X})\left[k(\mathbf{X}, \mathbf{X}) + \sigma_n^2\mathbf{I}\right]^{-1}\mathbf{y} \qquad (6)$$

and the (noise-free) predictive variance

$$
\begin{aligned}
\mathbb{V}[f(\mathbf{X}_*)] &= k(\mathbf{X}_*, \mathbf{X}_*) - k(\mathbf{X}_*, \mathbf{X}) \\
&\quad \left[k(\mathbf{X}, \mathbf{X}) + \sigma_n^2\mathbf{I}\right]^{-1} k(\mathbf{X}, \mathbf{X}_*), \quad (7)
\end{aligned}
$$

where $\mathbf{I}$ is the identity matrix and $k(\mathbf{X}, \mathbf{X})$ refers to the covariance matrix built by evaluating the covariance function $k(\cdot, \cdot)$ for all pairs of all row vectors $(\mathbf{x}_i, \mathbf{x}_j)$ of $\mathbf{X}$.

To sum up, Eq. (6) provides the predictive mean for the deformation cost when carrying out a movement along $\mathbf{x}^*$ and Eq. (7) provides the corresponding predictive variance.

4.3 Efficient Regression by Problem Decomposition

With the GP model explained above, we can make predictions for a set of trajectories deforming an object given training data obtained from the simulation. The key problem in practice, however, is that a substantial set of training data is required to obtain accurate predictions of the deformation cost. For the objects we experimented with, around 3000 training trajectories are needed. The GP framework, however, has a runtime that is cubic in the number of training examples so that the approach gets rather inefficient for more than 1000 training examples.

Therefore, we decompose the overall regression problem into a number of local ones. For a query trajectory $\mathbf{x}^*$, we determine its M closest neighbors from the training data under our distance function given in Eq. (5) as

$$\mathbf{X}'(\mathbf{x}^*) = [\mathbf{x}_1'; \ldots \mathbf{x}_M'] = \underset{[\mathbf{x}_1'; \ldots; \mathbf{x}_M']}{\operatorname{argmin}} \sum_{k=1}^{M} d(\mathbf{x}_k', \mathbf{x}^*). \qquad (8)$$

The M closest neighbors $\mathbf{X}'$ to the query trajectory $\mathbf{x}^*$ are the training data points that have the highest influence on the prediction of y^* in the GP framework. Considering only $\mathbf{X}'$ instead of $\mathbf{X}$ in the GP is equivalent to assuming that $k(\mathbf{x}^*, \mathbf{x}_i) = 0$ for all $\mathbf{x}_i$ that are not part of $\mathbf{X}'$. In our current implementation, we are able to get accurate predictions by setting $M = 50$. We experienced that the loss is negligible with respect to larger values of M, at least in all of our experiments. Determining the M closest neighbors to $\mathbf{x}^*$ can be computed efficiently using a KD-tree that is built once from the training data. Thus, queries can be obtained in logarithmic time in the number of training examples and the GP prediction does not depend on the size of the training set anymore but only on M.

4.4 Considering the Full Kinematic Chain for Estimating the Deformation Cost

The deformation simulation system considers the movement of the robot's endeffector along the described trajectory to compute the deformation cost. It does not consider the full configuration of the arm. This is clearly an approximation but it allows us to parametrize the regression problem with a low-dimensional input. Otherwise, the full configuration of the robot would need to be considered in the GP framework. With higher-dimensional inputs, a much larger number of training examples would be needed. To take into account the fact that not only the end-effector but also other body parts may deform an object, we sample multiple points along the kinematic chain of the robot. Then, we perform the estimation of the deformation cost for all sampled points along the kinematic chain and consider the maximum of the individual costs

$$C_{def} = \max_b \mathrm{GP}(\mathbf{x}^*(b), \mathbf{X}'(\mathbf{x}^*(b)), \mathbf{y}'(\mathbf{x}^*(b))), \qquad (9)$$

where b refers to the individual body parts and $\mathbf{x}^*(b)$ to the motion that the body parts carry out given the kinematic structure of the robot. Considering the maximum in Eq. (9) instead of, for example, the sum, generates more accurate predictions since the largest deformation forces are typically generated by one body part only.

5 Experimental Evaluation

5.1 Prediction of Deformation Costs

In this section, we evaluate our GP-based regression technique for predicting the deformation costs of robot trajectories. To show the effectiveness of the GP-based technique, we furthermore compare it to a nearest-neighbor prediction, which uses the average of the M nearest neighbors as an estimate. Our deformable object is a plush teddy bear for which we estimated the deformation parameters. To learn the deformation cost function of the teddy bear, we generated a set of samples by performing deformation simulations for different trajectory parameters. Since the simulation of sample trajectories is time-consuming, we restrict the manipulation movements to movements in the plane at different z-levels. Note that this can easily be generalized to arbitrary trajectories in 3D.

We consider 3 different data sets, which are D1 with 1,800 trajectory samples at $z = 0$, 20, and 40 cm, D2 with 1,400 trajectory samples at $z = 10$ and 30 cm, and D12 which is the combination of D1 and D2 with 3,200 trajectory samples. To evaluate the accuracy of the deformation cost prediction,

Table 1: Performance comparison for GP-based regression and nearest-neighbor approximation.

Dataset	RMSE			∅ time (ms)	
	NN	GPStd	GPOpt	GPStd	GPOpt
	leave-one-out				
D1	24.3	18.4	9.2	26.3	48.2
D2	19.5	27.0	5.8	19.3	42.9
D12	18.0	15.2	7.5	46.9	69.7
	cross-validation				
D1 on D2	26.9	22.5	17.8	19.4	42.1
D2 on D1	17.3	14.6	9.4	25.0	46.5

Figure 3: Comparison of the prediction performance for nearest-neighbor estimation and GP-Regression (leave-one-out cross-validation on D12).

Figure 4: Comparison of the prediction performance for nearest-neighbor estimation and GP-Regression(top: cross-validation D2 on D1, bottom: cross-validation D1 on D2).

we performed two different experiments, namely leave-one-out cross-validation for D1, D2, and D12 as well as cross-validation of D1 on D2 and vice versa. We compare the prediction results for the 50 nearest-neighbor prediction (NN), the prediction of a GP with standard hyperparameters (GP-Std), and the prediction of a GP with optimized hyperparameters (GPOpt). The results for the different data sets are summarized in Tab. 1. Whereas a visual comparison for the leave-one-out validation is shown in Fig. 3, the results for the cross-validation are depicted in Fig. 4.

5.2 Performance

In this section, we analyze the computational overhead introduced by our estimation of the deformation cost compared to a standard roadmap planner and compare it to a planner with integrated deformation simulation.

Generation of path examples The computation of the example trajectories is done offline in a preprocessing step. The simulation of one example trajectory takes on average 40 s for a trajectory length of approximately 1 m. The total simulation time for obtaining the 3,200 path examples we used in our evaluation was around 129,417 s (approx. 36 h).

Roadmap Computation The roadmap for a given static and non-deformable environment is computed in a preprocessing step. The main computational load comes from collision checks that need to be performed in order to determine edges that can be connected by collision-free paths. This is independent of our deformation cost estimation and takes for our test scenario with 1,000 sample configurations and 7,306 collision-free edges around 40 min. Evaluating the deformation cost for the 7,306 edges additionally takes 105 s (note
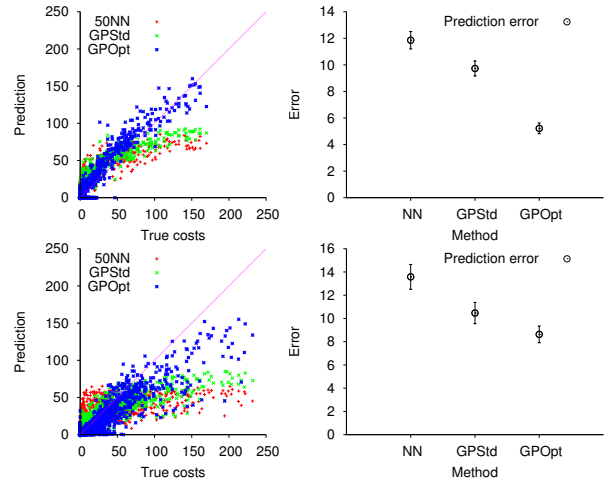
that only the edges intersecting the bounding sphere of the deformable object need to be further analyzed using our GP based regression. These were 801 edges in this example.

Answering Path Queries To answer path queries, starting and goal configurations need to be added to the roadmap. The planner attempts to connect these to the M nearest neighbors in the roadmap. The time-consuming factor here is the collision-checking. We evaluated 12 path queries. Connecting them to the roadmap took on average 3.5 s. The necessary evaluation of the deformation costs of the collision-free edges additionally requires 1.8 s on average.

Comparison to a Roadmap Planner with Integrated Deformation Simulation Instead of precomputing sample trajectories and estimating the deformation costs of edges in the roadmap using regression, it would be possible to perform the simulation of the edges directly when constructing the roadmap. Considering the example above, evaluating 801 edges requires an additional 267 min when constructing the roadmap. Furthermore, when answering path queries, we need to connect the start and the goal by adding new edges for which simulations need to be performed online. This requires another 10 min per path query. In contrast to that, our GP-based approach adds an overhead of approximately 1.8 s, thus requiring two orders of magnitude less computation time.

5.3 Example Trajectories

Finally, we carried out two planning experiments that are designed to illustrate the generated trajectories of our planner. We placed a deformable plush teddy bear in a shelf that is considered as a static obstacle. In both experiments, the robot had to move its arm from its current configuration to a goal. In the first setting, the target configuration was behind the teddy bear (Fig. 5) and in the second setting it was on the other side of the bear (Fig. 6). In both cases, the planner that does not consider the deformation costs would lead to
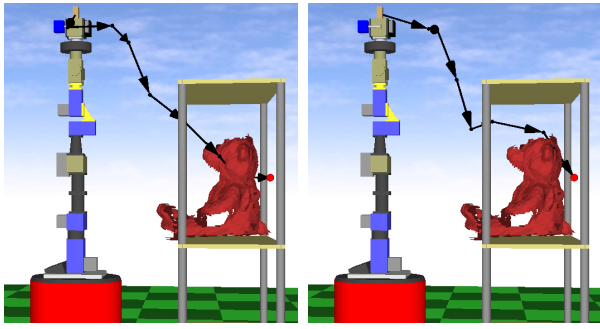
Figure 5: Planning example. Left image: shortest path, right image: trade-off between path cost and deformation cost.
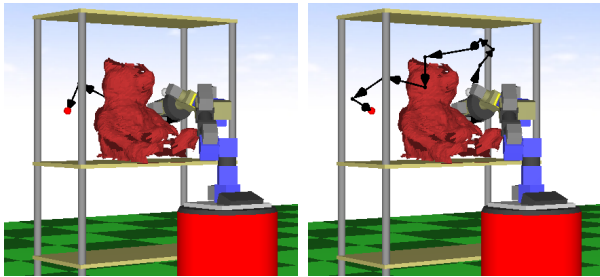

Figure 6: Planning example. Left image: shortest path, right image: trade-off between path cost and cost introduced by deforming the bear (only minimal deformations occur here).

significant deformation (left images) whereas our approach results in less deformation and still short paths (right images). Videos of the robot executing these trajectories in a deformation simulation can be found on our website[2].

6 Conclusion

In this paper, we presented a novel approach for efficiently planning the motion of a manipulation robot in environments that contain deformable objects. Our planner is based on probabilistic roadmaps and considers deformation costs that are computed from a physical simulation engine. To overcome the high computational demands of an appropriate physical deformation simulation, our approach employs an efficient variant of Gaussian process regression to estimate the deformation cost for individual objects based on training data. To limit the complexity of the regression problem, we train the Gaussian process only based on the most relevant training examples given a specific query trajectory. The training data is generated offline in a preprocessing step using a physical deformation simulation system. Consequently, no simulations are needed during runtime. Our experimental evaluation shows that our approach enables the robot to accurately estimate the expected deformation cost that its manipulator introduces to the objects in the scene along its path. It furthermore shows that our method substantially reduces the computation time compared to an approach that relies on the simulation engine during planning.

---

[2] http://www.informatik.uni-freiburg.de/~bfrank/defplan/

References

Anshelevich, E.; Owens, S.; Lamiraux, F.; and Kavraki, L. 2000. Deformable volumes in path planning applications. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Bayazit, O.; Lien, J.-M.; and Amato, N. 2002. Probabilistic roadmap motion planning for deformable objects. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Frank, B.; Becker, M.; Stachniss, C.; Teschner, M.; and Burgard, W. 2008. Efficient path planning for mobile robots in environments with deformable objects. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Frank, B.; Stachniss, C.; Schmedding, R.; Teschner, M.; and Burgard, W. 2009. Real-world robot navigation amongst deformable obstacles. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Frank, B.; Schmedding, R.; Stachniss, C.; Teschner, M.; and Burgard, W. 2010. Learning the elasticity parameters of deformable objects with a manipulation robot. In *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*.

Gayle, R.; Segars, P.; Lin, M.; and Manocha, D. 2005. Path planning for deformable robots in complex environments. In *Proc. of Robotics: Science and Systems (RSS)*.

Hauth, M., and Strasser, W. 2004. Corotational Simulation of Deformable Solids. In *Int. Conf. on Computer Graphics, Visualization, and Computer Vision (WSCG)*, 137–145.

Heidelberger, B.; Teschner, M.; Spillmann, J.; Mueller, M.; Gissler, M.; and Becker, M. 2006. DefColStudio 1.1.0. http://cg.informatik.uni-freiburg.de/software.htm.

Holleman, C.; Kavraki, L.; and Warren, J. 1998. Planning paths for a flexible surface patch. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Kavraki, L.; Svestka, P.; Latombe, J.-C.; and Overmars, M. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* 12(4):566–580.

Ko, J., and Fox., D. 2009. GP-BayesFilters: Bayesian filtering using gaussian process prediction and observation models. *Autonomous Robots*.

Maitin-Shepard, J.; M. Cusumano-Towner, J. L.; and Abbeel, P. 2010. Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Maris, B.; Botturi, D.; and Fiorini, P. 2010. Trajectory planning with task constraints in densely filled environments. In *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*.

Mueller, M., and Gross, M. 2004. Interactive Virtual Materials. In *Graphics Interface*, 239–246.

Rasmussen, C. E., and Williams, C. K. 2006. *Gaussian Processes for Machine Learning*. The MIT Press.

Rodríguez, S.; Lien, J.-M.; and Amato, N. 2006. Planning motion in completely deformable environments. In *Proc. of the Int. Conf. on Robotics & Automation (ICRA)*.

Vasudevan, S.; Ramos, F.; Nettleton, E.; and Durrant-Whyte, H. 2009. Gaussian process modeling of large scale terrain. *Journal of Field Robotics* 26(10).