

Efficient Path Planning for Mobile Robots in Environments with Deformable Objects

Barbara Frank

Markus Becker

Cyryll Stachniss

Wolfram Burgard

Matthias Teschner

Abstract—The ability to reliably navigate through the environment is an important prerequisite for truly autonomous robots. In this paper, we consider the problem of path planning in environments with non-rigid obstacles such as curtains or plants. We present an approach that combines probabilistic roadmaps with a physical simulation of object deformations to determine a path that optimizes the trade-off between the deformation cost and the distance to be traveled. We describe how our approach utilizes Finite Element theory for calculating the deformation cost. Since the high computational requirements of the corresponding simulation prevent this method from being applicable online, we present an approximation that uses a preprocessing step to determine a deformation cost function for each object. This cost function allows us to estimate the deformation costs of arbitrary paths through the objects and is used to evaluate the trajectories generated by the roadmap planner online. We present experiments which demonstrate that the resulting algorithm plans nearly identical trajectories compared to the method that relies on computationally intense simulations. At the same time, our approach allows the robot to quickly calculate paths in environments with deformable objects.

I. INTRODUCTION

Path planning is one of the fundamental problems in robotics, and the ability to plan collision-free paths is a precondition for numerous applications of autonomous robots. The path planning problem has traditionally received considerable attention in the past and has been well-studied. The majority of approaches, however, has focused on the problem of planning paths in static environments and with rigid obstacles [16, 7, 17]. In the real world, not all obstacles are rigid and considering this knowledge can enable a robot to accomplish navigation tasks that otherwise cannot be carried out. For example, in our everyday life we deal with many deformable objects such as plants, curtains, or cloth and we typically utilize the information about the deformability of objects when planning a movement. As a motivating example, consider the situation depicted in Fig. 1, in which the robot needs to pass through a curtain to move from its current position to the goal location since no other path exists in the environment. In this particular situation, traditional approaches that do not take the deformability of objects into account, will fail since no collision-free path exists. In contrast to this, the approach presented in this paper is able to determine the deformation cost introduced by passing the curtain and to utilize this information during path planning. The key idea of our approach is to use a heuristic




Fig. 1. Path planning in an environment with a deformable object. The robot (shown in red) deforms the curtain on its path to the goal.

function to estimate the deformation cost, which allows the robot to perform the necessary calculations online.

One potential method of taking deformations of objects into account is by generating trajectories using a method such as probabilistic roadmaps (PRM) [15] and considering deformable objects as free space. When answering path queries, the planner has to simulate the deformation of the non-rigid objects resulting from the interaction with the robot. However, performing an appropriate simulation typically requires significant computational efforts which makes such an approach unsuitable for online trajectory planning. Therefore, we propose an approach to learn an approximative deformation cost function in a preprocessing step. The advantage of our method is that this function can be evaluated efficiently during planning. In this way, our approach reduces the time to solve a path query from several minutes to a few hundred milliseconds. The assumption made throughout this paper is that the robot can deform but cannot move objects in the environment. Additionally, we neglect the interactions between different deformable objects.

The contribution of this paper is an approach to mobile robot path planning that explicitly considers deformable objects in the environment. It employs the probabilistic roadmap method and learns a deformation cost function using a physical simulation engine that is based on Finite Element theory. Our approach trades off the travel cost with the deformation cost when answering path queries and can be executed online.

This paper is organized as follows: After discussing related work, we present in Section III our technique to compute the deformations of objects by means of physical simulation. We then describe how to plan a path in the presence of deformable objects. Additionally, we describe our approximation of the deformation cost function in Section IV. Finally, we present experiments that illustrate the advantages of our approach compared to previous methods.

All authors are with the Department of Computer Science, University of Freiburg, 79110 Freiburg, Germany.

{bfrank,mbecker,stachnis,burgard,teschner}@informatik.uni-freiburg.de

II. RELATED WORK

The majority of approaches to mobile robot path planning assumes that the environment is static and that all objects are rigid [16, 15, 3]. Recently, several path planning techniques for deformable robots in static environments have been presented [4, 10, 14]. In the context of deformable objects, the underlying model has a substantial influence on the accuracy of the estimated deformations as well as on the performance of the planner. One typically distinguishes between geometrically and physically motivated approaches. Geometric approaches such as the free-form deformation (FFD) can be computed efficiently. For example, the FFD method of Sederberg and Parry [21] is based on trivariate Bernstein polynomials and allows for deformation by manipulating the control points.

To represent non-rigid objects and to calculate deformations, mass-spring systems have been frequently used. They are easy to implement and can be simulated efficiently. Whereas such models are able to handle large deformations, their major drawback is the tedious modeling as there is no intuitive relation between spring constants and physical material properties in general [19]. Finite Element Methods (FEMs) reflect physical properties of the objects in a better way. This allows for a more intuitive modeling since they require only a small number of parameters. The disadvantage of FEMs is the computational resources required to calculate deformations. In our current system we therefore use the computationally efficient co-rotational Finite Element approach [11, 18] which avoids nonlinear computations.

In the context of path planning, Kavraki *et al.* [14] developed the f-PRM-Framework that is able to plan paths for flexible robots of simple geometric shapes such as surface patches [13] or simple volumetric elements [2]. They apply a mass-spring model to compute deformations. The planner selects the deformation of the robot that minimizes its deformation energy. Similar to this technique, Gayle *et al.* [10] presented an approach to path planning for a deformable robot that is based on PRMs. To achieve a more realistic simulation of deformations they add constraints for volume preservation to the mass-spring model of the robot. Bayazit *et al.* [4] also studied planning for a deformable robot. Their algorithm proceeds in two steps. First, an approximate path is found in a probabilistic roadmap. In the next step, this path is refined by applying a free-form deformation to the robot and hence avoiding collisions with obstacles. This deformation method can be computed efficiently but is less accurate than physically motivated models. In contrast to our approach, these planners deform the robot rather than the obstacles to avoid collisions.

An approach to planning in completely deformable environments has been proposed by Rodríguez *et al.* [20]. They employ a mass-spring system with additional physical constraints for volume-preservation [23] 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. The




Fig. 2. Different levels of representation for a deformable object (left) in the simulation environment: fine surface mesh (middle) and coarse tetrahedral mesh (right).

obstacles in the environment are deformed through external forces resulting from collisions with the robot.

All techniques mentioned above, however, require substantial computational resources and cannot be executed online in general. In contrast to this, our approach can efficiently answer path queries by estimating potential deformations of objects in a preprocessing step. This is achieved by approximating a deformation cost function which is then considered during the planning process. Furthermore, our deformable model is based on FEM, which allows for more accurate deformations.

III. SIMULATION OF DEFORMABLE OBJECTS

To consider non-rigid obstacles in the environment during planning, we need a model that allows us to compute the deformations given an external force. In this section, we describe how we achieve a physically realistic simulation of object deformations. We will first introduce the co-rotational FEM. Then, we describe how to detect collisions between deformable objects and the robot and how to compute contact forces resulting from collisions. Finally, we introduce the cost resulting from a deformation.

Our simulation system proceeds as follows: in each time step, it computes deformations and unconstrained motions of objects, then it detects collisions, computes contact forces for colliding points, and finally corrects the unconstrained motion with appropriate repulsion forces.

A. Deformable Modeling

The obstacles in the workspace are 3D objects. The surface of objects is represented by a fine resolution triangle mesh. A tetrahedral mesh is used to represent the interior of these volumetric objects (see Fig. 2). The actual deformations are computed based on the coarse resolution tetrahedral mesh.

To compute the deformation of our tetrahedral objects we use the co-rotational Finite Element formulation [11, 18]. The total potential energy of a single tetrahedral element e is given by

$$\Pi = U_e + WP, \quad (1)$$

where WP is the work potential. WP is determined by the external forces and inner energy U_e

$$U_e = \frac{1}{2} \int_e \sigma^T \epsilon dV. \quad (2)$$

Since we assume only linear isotropic materials, we have a linear relation between the stress σ and the strain ϵ given by the generalized Hooke's law.

The key idea of the Finite Element method is to discretize the object into a finite set of elements (in our case tetrahedrons) to compute the deformations based on Eq. (1) on the nodes and to interpolate the deformation in the elements using the nodal values. To compute the strain ϵ from the nodal deformations in our model, we use the linear Cauchy strain tensor which is efficient to compute. The Cauchy tensor, however, is not rotationally invariant. This leads to ghost forces which result in distortions for large rotational deformations. Therefore, we keep track of the rigid body motion for each element by extracting the rotation from the transformation matrix using polar decomposition. Applying the strain tensor in the rotated frame, leads to rotational invariance and has low computational costs compared to the nonlinear strain tensor. We will discuss the performance of our Finite Element approach compared to the versatile mass-spring approach used by Rodríguez *et al.* [20] in the experimental results section.

B. Collision Detection

For the realistic processing of interactions between the robot and the deformable objects, an efficient collision detection algorithm is required. We employ a spatial subdivision scheme in our simulation system, where the elements are stored in a hash grid [24]. The key idea of this approach is to implicitly discretize \mathbb{R}^3 by storing the elements and nodes in the hash grid. Since the space is usually filled sparsely and non-uniformly, this method consumes less memory than an explicit discretization. The hash key is computed from the coordinates of the corresponding grid cell. As a result, only the elements with the same hash key need to be checked for collisions.

To check for collisions, one computes the intersection between points and tetrahedra. This can be done efficiently using barycentric coordinates of the points with respect to the tetrahedra.

Methods commonly employed for rigid bodies such as bounding box hierarchies [8] are less suited for deformable objects, since these spatial data structures cannot be precomputed [25].

C. Computation of Contact Forces

To handle collisions between the robot and the deformable objects, we employ the force-based collision handling scheme proposed by Spillmann *et al.* [22]. It combines the advantages of penalty and constraint-based collision handling schemes. For a set of colliding points of a tetrahedral mesh, we compute a collision free state using a linearized relation between internal forces and displacements of all affected points. Contact forces can be computed analytically to obtain this collision-free state while conserving overall system energy.

Using this combination of FEM-based simulation and the collision handling described above, our system can simulate thousands of tetrahedra at interactive rates. An example implementation of the simulation system is available online [12].

D. Object Deformation Costs

The inner energy of an object, specified in Eq. (2), provides a measure of the deformation costs of a tetrahedral object and thereby of the additional energy consumption of the robot. In case of an undeformed object, the inner energy is zero. Otherwise, the inner energy increases depending on the deformation of the object. For an object O with tetrahedral elements $\{e_i\}$, we define the total inner energy U_O induced by a robot r at position \mathbf{p} approaching from direction θ by the sum over the inner energies of all elements e_i of the object $U_O(r, \mathbf{p}, \theta) := \sum_{e_i \in O} U_{e_i}(r, \mathbf{p}, \theta)$.

For any given position \mathbf{p} and direction θ we determine the total deformation cost $C_{def}(\mathbf{p}, \theta) := \sum_O U_O(r, \mathbf{p}, \theta)$ by summing over all objects O in our workspace. The direction θ has to be taken into account, since deformable objects might deform differently when approaching them from different directions.

The total deformation cost on a path Γ of the robot approaching from direction θ in our environment naturally results in the sum of the deformation costs of all objects that are deformed by the robot while moving on the path in discrete time steps t_i :

$$C_{def}(\Gamma, \theta) = \sum_{t=t_1}^{t_n} C_{def}(\mathbf{p}_r(t_i), \theta(t_i)). \quad (3)$$

Here $\mathbf{p}_r(t_i)$ is the position of the robot on Γ at time t_i . $\theta(t_1)$ is given by θ , all other directions $\theta(t_i)$ are determined as the difference between $\mathbf{p}_r(t_i)$ and $\mathbf{p}_r(t_{i-1})$.

IV. PATH PLANNING WITH DEFORMABLE OBJECTS

In this section, we introduce an approximative deformation cost function that allows a robot to plan a path in such environments online.

A. Overview of the Path Planning System

The general path planning problem is to find a sequence of valid robot configurations that lead from the starting to the goal configuration. Probabilistic roadmap planners achieve this by constructing a roadmap that represents the environment of the robot and by applying a graph search algorithm to find a path from the starting to the goal configuration. The roadmap is constructed in a preprocessing step by sampling points in the configuration space of the robot. These points have to satisfy certain feasibility constraints. In general, valid configurations are required to be part of the free configuration space \mathcal{C}_{free} . In our situation, however, we modify this constraint and require configurations to be in $\mathcal{C}_{free} \cup \mathcal{C}_{def}$. Thus, we also accept configurations that lead to collisions with deformable objects. In our current implementation, we use Hammersley-sampling [6] to generate configuration hypotheses in the space. This deterministic sampling method generates a sequence of points that are distributed with low discrepancy. After a designated number of samples has been generated, a local planner connects neighboring samples




Fig. 3. Deformation costs for moving a robot along straight lines through a curtain. The lines are specified by starting points (x, y) and travel direction $\theta = 0^\circ$ relative to the center of mass of the obstacle.

for which a valid path exists. This typically results in a roadmap that covers the environment of the robot and can be utilized for planning paths on which objects are allowed to be deformed by the robot. To answer a path query, we then insert the starting and the goal configuration into the roadmap and connect them to their neighbors. Finally, we apply A^* to find the best path from the starting to the goal location on the graph. Here, we search for the path with the best trade-off between travel costs and deformation cost. Therefore, we need to be able to estimate the expected deformation cost arising on the edges of the roadmap.

The simulation system presented in the previous section can be used inside the planning algorithm to compute the deformation cost $C_{def}(i)$ of an edge i by simulating a robot moving along this edge deforming the object. The edges considered during A^* planning are evaluated by trading off the deformation and travel cost. In our planning system, we assume the travel cost to be proportional to the length of the edge i . This results in the cost function

$$C(i) := \alpha C_{def}(i) + (1 - \alpha) \text{length}(i), \quad (4)$$

where $\alpha \in [0, 1]$ is a user-defined weighting coefficient. In order to obtain an admissible heuristic for A^* , we use the Euclidean distance to the goal location 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 for a given user-defined parameter α .

This technique leads to a working planning system that considers deformations of the objects in the environment when planning a trajectory for a mobile robot. The drawback of this technique, however, is its high computational requirements. Answering a path query typically takes several minutes even for small examples. Therefore, we developed an alternative approach that computes an approximation of the cost function in advance and thus facilitates online path planning.

B. Approximative Deformation Cost Function

The goal of the approximative cost function is to quickly provide an estimate of the deformation costs for all objects along an edge in the roadmap. Such a function can be used in the planning approach described above to determine C_{def} . The actual value of the deformation cost function of an object mainly depends on the trajectory of the robot relative to the object. Therefore, we precompute the deformation cost for a number of line segments through each object.




Fig. 4. Test environments: world 1 with curtains (left), world 2 with rubber ducks (middle), and world 3 with rubber ducks and curtains (right).

A line is specified by a starting location (x, y) and the traveling direction θ as well as the traveled distance l on the line segment. The traveled distance is constrained to the maximum distance that the robot can move while still deforming the object.

In a preprocessing step, we carry out the simulations for a uniform resolution of starting points and directions and store the deformation costs for a fine length resolution in a histogram. This leads to the approximate deformation cost function $\hat{C}_{def}(x, y, \theta, l) \rightarrow \mathbb{R}$ which returns the deformation cost for edges of the roadmap.

We compute the deformation cost $\hat{C}_{def}(x, y, \theta, l)$ of an arbitrary edge e in the roadmap by first determining the starting position (x, y) , direction θ , and length l relative to the deformable object. We then apply a kernel smoother [1] considering all neighboring line segments e^t in the histogram

$$\hat{C}_{def}(e) = \frac{\sum_t K\left(\frac{e-e^t}{h}\right) \hat{C}_{def}(e^t)}{\sum_t K\left(\frac{e-e^t}{h}\right)}, \quad (5)$$

where $K(u)$ is the multivariate Gaussian kernel with identity as covariance.

As distance metric between different line segments, we employ the Euclidean distance and normalize the orientation with respect to the positions.

To finally answer path queries, we apply the A^* algorithm on the roadmap. The cost of each edge in the graph is computed according to Eq. (4) using the precomputed approximation \hat{C}_{def} of C_{def} . Computing the deformation costs in a preprocessing step substantially increases the performance of our planner as can be seen in the experiments.

Although the precomputation is computationally intense, it has to be done only once for each distinct object. Such a cost function can even be transferred between environments. The following section provides results on the runtime of the precomputation for different resolutions of the deformation cost grid. Fig. 3 illustrates the deformation cost \hat{C}_{def} of the curtain shown in Fig. 1 along a series of straight lines.

V. EXPERIMENTS

We carried out different experiments to evaluate our path planning approach. In this section, we first compare the deformation cost obtained by the FEM-based simulation technique with our approximative solution that computes a deformation cost function for each object in a preprocessing step. Next, we investigate how the deformation cost weighting coefficient α influences the path search. Finally, we present examples of planned trajectories in environments with deformable objects.




Fig. 5. Comparison of the error in the deformation cost for the simulation-based approach and our approximation.




Fig. 6. Example trajectories generated by the simulation approach (left) vs. our approximation approach (right).

A. Cost Function and Runtime

In the first experiment, we compare the simulation approach to estimate the deformation cost with our approximative cost function. We chose curtains and rubberducks as deformable objects. The curtains are modeled to be easily deformable while the rubberducks have high deformation costs. Both approaches had to solve 25 path queries in the three test environments depicted in Fig. 4. After planning, the best trajectory is sent to a path execution module that guides the robot along that trajectory. In our simulation, the execution of motion commands is affected by noise.

The experiments are carried out for different resolutions of the approximate cost function. Furthermore, we compare our approach to a simulation system using the versatile mass-spring model. We evaluated the error between predicted and measured deformation cost. The results are shown in Fig. 5. As expected, the error of the simulation technique is typically smaller compared to our approximative approach. This, however, comes at the expense of running time as illustrated in Table I. While our approach answers path queries even for complex environments in a few hundred milliseconds on average, the simulation approach spends generally about half an hour on a single query. Thus, our approach is about four orders of magnitude faster than the simulation using FEM. We also compared our approach to a simulation system using the versatile mass-spring model. Although this simulation system can be evaluated faster, our approach still is about 2000 times faster. The runtime for the precomputation of the approximate cost function for different resolutions of the cost grid is summarized in Table II.

Additionally, we carried out an experiment in a randomized world, where we compare the computed paths for the simulation and the approximation approach. The

TABLE I

AVERAGE RUNTIME INCLUDING CONFIDENCE INTERVALS.

World	\varnothing Query (FEM)	\varnothing Query (mass-spring)	\varnothing Query (our approach)
1	36 m 41 s \pm 347 s	12 m 45 s \pm 112 s	0.4 s \pm 0.04 s
2	30 m 33 s \pm 512 s	8 m 10 s \pm 73 s	0.2 s \pm 0.02 s
3	29 m 43 s \pm 130 s	7 m 27 s \pm 36 s	0.3 s \pm 0.04 s

TABLE II

RUNTIME FOR THE PRECOMPUTATION OF THE DEFORMATION COST FUNCTION FOR DIFFERENT RESOLUTIONS OF THE COST GRID.

Object	Mass-spring simulation		FEM simulation	
	coarse res. (200 lines)	fine res. (7056 lines)	coarse res. (200 lines)	fine res. (7056 lines)
curtain	17 m 48 s	10 h 37 m	1 h 11 m	41 h 16 m
rubberduck	18 m 21 s	10 h 56 m	1 h 16 m	44 h 44 m

generated trajectory points deviate on average by 0.09 m in an environment of 2.6×9 m and the deformation costs of the trajectories deviate by $9.4 \pm 5.2\%$.

In most cases, the actual trajectories reported by the different planners do not deviate substantially. As the examples depicted in Fig. 6 illustrate, the resulting trajectories are similar. This suggests that our approximative solution provides appropriate trajectories for planning in environments with deformable objects.

B. Determining the Weighting Coefficient

Eq. (4) contains the weighting factor α that trades off the travel costs with deformation costs. To find good values for this factor, we carried out a series of planning experiments with varying values for α . Low values for α result in the fact that the robot traverses objects that are hard to deform in order to obtain a short trajectory. In contrast to this, high values for α will lead to a planning system that




Fig. 7. Different trajectories obtained in two environments depending on the weighting coefficient α .




Fig. 8. Deformation and travel costs of executed trajectories depending on the weighting coefficient α . The left plot corresponds to the environment in the first row of Fig. 7 and the right one to the one in the second row.




Fig. 9. Example trajectory guiding the robot through a deforming object.

entirely avoids deformations if possible. Examples for such trajectories are depicted in Fig. 7. The corresponding analysis of the deformation and travel costs is shown in Fig. 8. Based on these experiments, we generally set $\alpha = 0.2$. As a result, the robot selects trajectories through easily deformable objects such as curtains and tries to avoid objects that cause high deformation costs such as the rubber ducks.

Finally, Fig. 9 shows a sequence of snapshots taken during a planning experiment. They illustrate that a robot using our planning approach selects trajectories through deformable objects in case the deformation is not too expensive. More examples and animations are available at our web site [9].

VI. CONCLUSIONS

In this paper, we presented an approach to path planning in environments with non-rigid objects. Our planner takes potential deformations of objects into account using a simulation engine that is based on the physically accurate Finite Element method. To improve the efficiency of the planner, our approach uses a pre-computed and object dependent deformation cost function that estimates the deformation cost of path segments relative to the object. The cost function is learned offline and is integrated into a probabilistic roadmap planner. To calculate paths, our system trades off deformation and travel costs. As a result, we obtain highly efficient paths and at the same time avoid computationally expensive simulations during runtime.

Our approach has been implemented and tested exhaustively in environments with deformable objects. The utilization of the approximative cost function leads to a speedup of about four orders of magnitude compared to a system that performs the simulations at runtime.

Despite these encouraging results, we envision several aspects for further improvements. One of our next goals is to acquire models of real obstacles with our robot and estimate their elasto-mechanical parameters, for example by using the method proposed in our previous work [5]. This will allow for applying our system to real world settings and accurately considering the properties of real deformable objects.

VII. ACKNOWLEDGMENTS

This work has partly been supported by the German Research Foundation (DFG) under contract number SFB/TR-8.

REFERENCES

- [1] E. Alpaydin. *Introduction to Machine Learning*. MIT Press, 2004.
- [2] E. Anshelevich, S. Owens, F. Lamiroux, and L.E. Kavraki. Deformable volumes in path planning applications. In *Proc. of the IEEE Int. Conf. on Robotics & Automation*, pages 2290–2295, 2000.
- [3] J. Barraquand, B. Langlois, and J.-C. Latombe. Numerical potential field techniques for robot path planning. *IEEE Transactions on Systems, Man and Cybernetics*, 22(2):224–241, 1992.
- [4] O.B. Bayazit, J.-M. Lien, and N.M. Amato. Probabilistic roadmap motion planning for deformable objects. In *Proc. of the IEEE Int. Conf. on Robotics & Automation*, pages 2126–2133, 2002.
- [5] M. Becker and M. Teschner. Robust and efficient estimation of elasticity parameters using the linear finite element method. In *Proc. of Simulation and Visualization*, pages 15–28, 2007.
- [6] M.S. Branicky, S.M. LaValle, K. Olson, and L. Yang. Quasirandomized path planning. In *Proc. of the IEEE Int. Conf. on Robotics & Automation*, pages 1481–1487, 2001.
- [7] H. Choset, K.M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L.E. Kavraki, and S. Thrun. *Principles of Robot Motion*. MIT Press, 2005.
- [8] C. Ericson. *Real-Time Collision Detection*. Morgan Kaufmann, 2005.
- [9] B. Frank. Motion planning with deformable objects, 2008. <http://www.informatik.uni-freiburg.de/~bfrank/defplan>.
- [10] R. Gayle, P. Segars, M.C. Lin, and D. Manocha. Path planning for deformable robots in complex environments. In *Proc. of Robotics: Science and Systems (RSS)*, pages 225–232, 2005.
- [11] M. Hauth and W. Strasser. Corotational Simulation of Deformable Solids. In *WSCG*, pages 137–145, 2004.
- [12] B. Heidelberger, M. Teschner, J. Spillmann, M. Mueller, M. Gissler, and M. Becker. DefColStudio – interactive deformable modeling framework. <http://cg.informatik.uni-freiburg.de/software.htm>.
- [13] C. Holleman, L.E. Kavraki, and J. Warren. Planning paths for a flexible surface patch. In *Proc. of the IEEE Int. Conf. on Robotics & Automation*, pages 21–26, 1998.
- [14] L.E. Kavraki, F. Lamiroux, and C. Holleman. Towards planning for elastic objects. In *Robotics: The Algorithmic Perspective*, pages 313–325. A.K. Peters, 1998. Proc. of the Third Workshop on the Algorithmic Foundations of Robotics (WAFR).
- [15] L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [16] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Pub., 1991.
- [17] S.M. LaValle. *Planning Algorithms*. Cambridge Univ. Press, 2006.
- [18] M. Mueller and M. Gross. Interactive Virtual Materials. In *Graphics Interface*, pages 239–246, 2004.
- [19] A. Nealen, M. Mueller, R. Keiser, E. Boxerman, and M. Carlson. Physically Based Deformable Models in Computer Graphics. *Computer Graphics Forum*, 25(4):809–836, 2006.
- [20] S. Rodríguez, J.-M. Lien, and N.M. Amato. Planning motion in completely deformable environments. In *Proc. of the IEEE Int. Conf. on Robotics & Automation*, pages 2466–2471, 2006.
- [21] T.W. Sederberg and S.R. Parry. Free-form deformation of solid geometric models. In *Proc. of the Conf. on Computer graphics and interactive techniques*, pages 151–160, 1986.
- [22] J. Spillmann, M. Becker, and M. Teschner. Non-iterative computation of contact forces for deformable objects. *Journal of WSCG*, 15(1–3):33–40, 2007.
- [23] M. Teschner, B. Heidelberger, M. Mueller, and M. Gross. A versatile and robust model for geometrically complex deformable solids. In *Proc. of Computer Graphics International*, pages 312–319, 2004.
- [24] M. Teschner, B. Heidelberger, M. Mueller, D. Pomeranets, and M. Gross. Optimized spatial hashing for collision detection of deformable objects. In *Proc. Vision, Modeling, Visualization (VMV)*, pages 47–54, 2003.
- [25] M. Teschner, S. Kimmeler, B. Heidelberger, G. Zachmann, L. Raghupathi, A. Fuhrmann, M.P. Cani, F. Faure, N. Magnenat-Thalmann, and W. Strasser. Collision Detection for Deformable Objects. *Computer Graphics Forum*, 24(1):61–81, 2005.