

Gaussian Process Implicit Surfaces for Shape Estimation and Grasping

Stanimir Dragiev

Marc Toussaint

Michael Gienger

Abstract—The choice of an adequate object shape representation is critical for efficient grasping and robot manipulation. A good representation has to account for two requirements: it should allow uncertain sensory fusion in a probabilistic way and it should serve as a basis for efficient grasp and motion generation. We consider Gaussian process implicit surface potentials as object shape representations. Sensory observations condition the Gaussian process such that its posterior mean defines an implicit surface which becomes an estimate of the object shape. Uncertain visual, haptic and laser data can equally be fused in the same Gaussian process shape estimate. The resulting implicit surface potential can then be used directly as a basis for a reach and grasp controller, serving as an attractor for the grasp end-effectors and steering the orientation of contact points. Our proposed controller results in a smooth reach and grasp trajectory without strict separation of phases. We validate the shape estimation using Gaussian processes in a simulation on randomly sampled shapes and the grasp controller on a real robot with 7DoF arm and 7DoF hand.

I. INTRODUCTION

Grasping is a major challenge in robotics research. How a target object can be grasped is constrained to a large extent by its shape. The choice of an adequate object shape representation is therefore critical for fluent grasp motion control.

Many possible ways and levels to represent knowledge about objects of interest have been proposed, including direct visual cues [12], superquadrics – as primitives or in unions – to approximate shapes [19], voxel and octree based shape approximations [8] and standard mesh representations. What is a good choice to represent information of object shapes for grasping? Our view is that a shape representation has to account for two requirements: first, the representation has to serve directly as a basis to formulate a motion controller to generate a fluent reach and grasp motion. Second, the shape representation needs to fuse the available uncertain sensor information and thereby yield sufficient information for action selection and control.

Let us express these requirements more clearly. Information about the shape of an object can be acquired via visual sensors, via laser sensors, or – if in direct contact – via haptic sensors. Ideally, an internal object shape representation should provide a means to fuse – typically in the Bayesian sense – all these uncertain sensor channels

to a shape estimate of the object. For instance, when bad lighting conditions allow for only poor image quality, we should be able to represent a very uncertain shape estimate; when the object gets in reach we get additional more precise haptic feedback and should be able to combine both sensory sources. However, the need of sensor fusion for shape estimation is only half the story. Ideally, an internal object shape representation should also provide a convenient basis for manipulation, especially reach and grasp motion generation. In other terms, accurate shape estimation is not the primary objective; rather it is the combination of all sensorial information to support the control of a reach and grasp motion.

Given this motivation, we propose an object shape representation which allows to fuse sensorial information in a probabilistic manner – and therefore to represent shape uncertainty – and which can be used directly by a motion controller for fluent grasping movement generation. Our approach is based on a Gaussian process implicit surface representation.

Using implicit surfaces to represent shapes has a long tradition starting with a method to visualize 3D models of molecules in the early 80s [2]. Ever since they reached high esteem through variety of analytical algorithms for deformation and blending, e.g. [3]. Also in the Machine Learning community, implicit surface shape representations have been investigated. Methods include the Multi-level partition of unity implicits by Ohtake et al. [10] for local approximation; Curless and Levoy’s approach to integration of multiple range images [4]; an approach for approximation and deformation of 3D shapes with Support Vector Machines by Steinke et al. [14].

Since probabilistic sensor fusion implies the necessity to also represent uncertain shapes we need to extend implicit surfaces to uncertainty. Gaussian processes [11] (GP) are a standard method in Machine Learning to represent uncertainty – we therefore propose to use Gaussian processes as surface potential, such that the uncertainty in function value implies the shape uncertainty of our estimate. In a related work Williams and Fitzgibbon [18] employ a GP in a setting similar to [14] and show comparable results with the additional benefit of GP’s probabilistic interpretation. They derive a GP with thin plate covariance function from the thin plate regularizer and argue towards some virtues of their covariance vs. squared exponential. We will explain in detail in section II how sensor information (from vision, laser, or tactile) can be translated to conditioning a Gaussian process such that the implied implicit surface corresponds to the fused shape estimate from this sensor information. Here

Stanimir Dragiev and Marc Toussaint are with the Machine Learning group at the TU Berlin, Germany
stanio, mtoussai@cs.tu-berlin.de

Michael Gienger is with Honda Research Institute Europe, Offenbach, Germany
michael.gienger@honda-ri.de

Stanimir Dragiev gratefully acknowledges the financial support from Honda Research Institute Europe.

should be mentioned that both [18] and [14] use surface points and off-surface points for training the regression.

The mean of the posterior of the GP, which we refer to as implicit surface potential (ISP), is suited as a basis for a smooth reach and grasp controller. Most existing literature on grasp optimization focuses on the grasp itself, isolated from the reaching movement. For instance, [15] reviews the various literature on defining grasp quality measures, [13] learn which grasp positions are feasible for several objects, [6] efficiently compute good grasps depending on how the objects shall be manipulated, and [9] simplify the grasp computation based on abstracting objects into shape primitives. The coupling to the problem of reaching motion optimization is rarely addressed. A recent approach [1] makes a step towards solving the coupled problem by including a “environment clearance score” in the grasp evaluation measure. In that way, grasps are preferred which are not prohibited by immediate obstacles directly opposing the grasp. In [17] a method for combined reach and grasp optimization for a known object shape was presented. Our approach in this paper is to use the implicit surface potential that was learnt to estimate the object shape directly to also control reaching and grasping. The controller relies on the potential function steering the orientation of the fingers and wrist and eventually results in a smooth trajectory and feasible grasp posture. This use of potential function can to some degree be compared to classical approaches on using potentials for navigation and obstacle avoidance in mobile robots, e.g. by [7].

In the following we introduce briefly implicit surface and Gaussian process and describe the particular class of implicit surface we use. Thereafter we explain the use of GP for estimating ISF from sensor observations. Section III introduces several heuristics for achieving a feasible grasp based on a potential function. Two experiments for evaluating the methods are presented in section IV, followed by the conclusions.

II. GAUSSIAN PROCESSES FOR IMPLICIT SURFACE ESTIMATION

In the following we describe how to use Gaussian processes to represent uncertain shape estimates.

A. Challenges for GP shape estimation

Concerning the shape estimation part, we aim at integrating different sensor inputs, in particular the methods have to be able to deal with haptic, visual and laser feedback. Visual and laser sensors provide uncertain 3D coordinates of points on the surface of possibly far objects, with laser typically being more precise than 3D vision (based on stereo triangulation of key points). Haptic sensing can be more precise than vision, but is useful only with objects which are in reach. In addition to the position of a surface point, haptic feedback provides also an estimate of the local tangent plane of the shape at the point of contact. The above three examples illustrate the need to have a representation that




Fig. 1. Objects represented by an implicit surface function. Left: 1D object in $[-2, 2]$; right: 2D object described by the blue ellipse in the center of the plane.

allows to fuse uncertain multi-modal sensor information in a probabilistic shape estimate.

B. Implicit surface

We propose to describe an object by means of an implicit surface which in turn is described by Gaussian process. An implicit surface describes the shape of an object by means of a function which tells for each location in space whether it is part of the object or not. More formally, we define an implicit surface as the 0-level set of a real-valued potential function f with

$$f : \mathbb{R}^d \rightarrow \mathbb{R}; f(\mathbf{x}) \begin{cases} = 0, & \mathbf{x} \text{ on the surface} \\ > 0, & \mathbf{x} \text{ outside the object} \\ < 0, & \mathbf{x} \text{ inside the object} \end{cases} \quad (1)$$

In the following we call f the *implicit shape potential (ISP)*. Figure 1 gives two examples of objects represented by the 0-level set. Note that there are infinitely many potentials which result in the same shape. Generally, we will disambiguate this invariance by fixing a specific bias, i.e., value of the potential far from the object.

Further, we define the *implicit shape field (ISF)* as the negative gradient $-\nabla f(\mathbf{x})$ of the implicit shape potential $f(\mathbf{x})$. It points towards the object and will play a crucial role in the design of the reach and grasp motion controller based on the ISP.

C. Gaussian process ISP

How can we condition a Gaussian process on sensor information such that the respective ISP represents an estimate of the object shape? Consider a minimalist 1D example: In the 1D world an object is a segment and its surface consists of the two points which delimit the segment. On Figure 2 the GP prior has a bias $\mu = 1$, and two observations at $x = 2$ and $x = 3$ condition the GP such that the 1D object location is estimated to be the segment $[2, 3]$. In the following we show what information do these observations carry.

1) *Surface point observations*: In the simplest case the sensors return a set S of 3D points which the sensor believes to be points on the object surface (e.g., laser points, stereo triangulated points, or contact points). With respect to the ISP this data implies that the value of the potential at these points is zero. Therefore, every point in the data set implies a *value-zero-observation* for the Gaussian process ISP. The initial bias of the GP is set positive to reflect the prior




Fig. 2. Learning an implicit surface representation for an object in [2, 3]. Initially $\mu(\mathbf{x}) = 1$. Belief after one observation of surface point and normal (left) and after two observations (right). Note that the variance of the posterior is a quantification of how precise and trustful the estimate is, given the present data and prior knowledge.




Fig. 3. Example of a tactile sensor touching an object. Relationship between the normals of the tangent planes of the object sensor at the contact point. $-\mathbf{n}_s(\mathbf{x}) = \mathbf{n}_o(\mathbf{x})$ (depicted as \leftarrow)

that empty space is more likely. The value-zero-observations then condition the posterior Gaussian process such that the respective implicit surface becomes non-empty and represent the posterior shape estimate.

2) *Surface normal observations*: Haptic sensors do not only give information on the location of a 3D surface point but also on the normal vector $\mathbf{n}_o(\mathbf{x})$ of the local surface tangent, as illustrated in Figure 3. Similar normal information can also be derived (with more uncertainty) from the direction of the visual contact or the laser ray. By definition of implicit surface, the normal vector of the tangent plane needs to be equal to the gradient direction of the implicit surface potential, i.e., $\nabla f(\mathbf{x}) = \mathbf{n}_o(\mathbf{x})$.

A feature of Gaussian processes is the ability to condition them also on gradient-observations, i.e., on knowing the gradient of the function at a certain location instead of only the value, see [11, p.191]. Therefore, a sensor data-point conditions the GP estimation in a two-fold manner: by implying a value-zero-observation as well as a gradient-observation.

For completeness we give here the specific equations necessary to incorporate the gradient observations into the GP. The predictive distribution of a Gaussian process is defined according to [11] by the following mean and variance:

$$\bar{f}_* = \mathbf{k}_*^T (K + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \quad (2)$$

$$\mathbb{V}[f_*] = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (K + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_* \quad (3)$$

where k is the covariance function, \mathbf{k}_* is the vector of covariances between the test point and all observation points,




Fig. 4. Random samples from GP with 2D domain. Sampled functions are produced with varying covariance width, σ_w^2 , rises from left to right in both series. Functions are cut through the 0-level set and positive values are light colored, negative – dark.

K is the Gram matrix and \mathbf{y} denotes the observed values. In this work we generally use a Gaussian covariance function $k_G(\mathbf{x}_i, \mathbf{x}_j) = \sigma_p^2 \exp(-\frac{(\mathbf{x}_i - \mathbf{x}_j)^T (\mathbf{x}_i - \mathbf{x}_j)}{2\sigma_w^2}) + \sigma_n^2 \delta_{ij}$. To incorporate the gradient observations in the covariance vector \mathbf{k} and the Gram matrix \mathbf{K} we need the derivative of the covariance function [11]. Using $\mathbf{x}_i, \mathbf{x}_j$ as the observed points; d_j as the index of the component in which the derivative $\frac{\partial f(\mathbf{x}_j)}{\partial x_{d_j}}$ is observed, δ_{ij} being Kronecker's delta, we get:

$$k\left(f(\mathbf{x}_i), \frac{\partial f(\mathbf{x}_j)}{\partial x_{d_j}}\right) = \frac{\partial k(\mathbf{x}_i, \mathbf{x}_j)}{\partial x_{d_j}} = \frac{(\mathbf{x}_i - \mathbf{x}_j) k_G(\mathbf{x}_i, \mathbf{x}_j)}{\sigma_w^2},$$

and

$$k\left(\frac{\partial f(\mathbf{x}_i)}{\partial x_{d_i}}, \frac{\partial f(\mathbf{x}_j)}{\partial x_{d_j}}\right) = \frac{\partial^2 k(\mathbf{x}_i, \mathbf{x}_j)}{\partial x_{d_i} \partial x_{d_j}} = \frac{(\sigma_w^2 \delta_{ij} - (\mathbf{x}_{id_i} - \mathbf{x}_{jd_i})(\mathbf{x}_{id_j} - \mathbf{x}_{jd_j})) k_G(\mathbf{x}_i, \mathbf{x}_j)}{\sigma_w^2}.$$

D. The resulting implicit shape prior

Our Gaussian process prior implies directly a prior over shapes. This prior depends on the choice of hyper parameters for the covariance function and bias. Since the definition of an implicit surface is invariant to scaling of function values we can – without loss of generality – fix the GP bias to $\mu = 1$. The resulting shape prior is then influenced by the covariance width σ_w^2 . Figure 4 illustrates in 2D the shape priors (and also the implicit prior of how shapes are distributed in space) depending on σ_w^2 . It shows a series of randomly sampled functions from a two-dimensional GP. They are produced by sampling random observation and using the resulting predictive mean. Different colors indicate positive and negative function values.

The observation precisions σ_n^2 reflect the precision of the used sensors. The GP allows to have precision associated with each observation. Furthermore, as commented in [10], sensors may have different precision along different axes. In our case this translates to a GP with multivariate noise in the input dimensions. See [5] for details on how input noise translates into posterior variance.

III. IMPLICIT SURFACE GRASPING

In this section we describe a feedback reach and grasp controller which uses the Gaussian process potential described in the previous section. It generates fluent trajectories without strictly separating reach from grasp motions. This is achieved by introducing several task space objectives and controlling their relative importance based on the information carried by the GP ISP and its associated gradient field. In other words, the objectives are active throughout the whole trajectory, only their importance varies over time according to ISP.

We follow a pure geometric approach to grasping, i.e. in the scope of the controller are situations in which merely the shape of the object has influence on the grasp. Other features of the object, like mass inhomogeneity, deformability, surface friction, etc. are neglected.

A. General approach

The implicit surface potential $f(x)$ and its field $-\nabla f(x)$ provide two essential pieces of information to guide a reach and grasp motion: $-\nabla f(x)$ indicates the direction to the object surface for every point in space while $f(x)$ indicates a proximity to the object: it approaches zero in the vicinity of the object and takes a value of 1 far from the object. Combining these two pieces of information we can define desired infinitesimal motions for the hand, fingers, and finger orientations. Roughly, the controller can be summarized as follows:

- Far from the object ($f(x) \approx 1$) the robot should direct the hand towards the object, move the hand towards the object, and open the hand.
- In the vicinity of the object ($0 < f(x) < 1$) the robot should move the fingers onto the surface and align the fingers with the surface tangent.
- At the surface ($f(x) \approx 0$) the robot should close the fingers until pressure sensors indicate the desired contact.

Using $f(x)$ to blend between these task settings leads to a smooth grasp motion without ad hoc separation in strictly isolated phases. We will realize this strategy by defining respective *task variables* (also called endeffector variables) for the hand and fingers and compute the motion as the optimum of a squared cost minimization problem – a straightforward extension of operational space control for multiple (regularized) task variables. For completeness we briefly define this optimization problem in the following and then define in detail the task variables and overall controller.

B. Control under multiple task variables

Let $q \in \mathbb{R}^n$ be the n -dimensional vector describing the current joint angles of the robot. We formulate the control problem in the dynamic domain, based on the current state $(q, \dot{q}) \in \mathbb{R}^{2n}$.

We assume we have m different *task variables* $y_i \in \mathbb{R}^{d_i}$ (with the i task variable being d_i -dimensional) and we are given the respective kinematic function $y_i = \phi_i(q)$ to compute the joint angles and its Jacobian $J_i = \frac{\partial \phi_i}{\partial q}$ at the current q . For each task variable we assume to have current

desired values y_i^* , desired velocities \dot{y}_i^* as well as associated precisions¹ ρ_i and ν_i for values and velocities, respectively. These task variables define a cost function:

$$c(q_t, \dot{q}_t) = \sum_{i=1}^m \rho_{i,t} [y_{i,t}^* - \phi_i(q_t)]^2 + \nu_{i,t} [\dot{y}_{i,t}^* - J_i \dot{q}_t]^2$$

Using a local linearization of $\phi_i(q_t)$ this can be expressed as a quadratic form in (q, \dot{q}) .

Apart from this task cost term we have a control cost term penalizing accelerations, proportional to $a^\top H a$ with $a = \dot{q}_{t+1} - \dot{q}_t$. (Operating directly on accelerations is a simplification from full dynamic operational space control.) The sum of the task and control cost terms is a quadratic term in the new state. At each time step the controller computes an acceleration that minimizes these quadratic costs. See [16] for more details on the dynamic controller with multiple task variables.

C. Task Variables and the Reach and Grasp Controller

The feedback controller is motivated by the Schunk-arm with attached 3-humanoid-finger hand available at our lab. It is based on the local ISP information and simple heuristics.

We define one task variable (TV) for wrist orientation, one for wrist position, one TV per finger for tip orientation, and one TV per finger for tip position. In addition, three feasibility TVs implement collision and joint limits avoidance. A TV for the skin sensory response is used to eventually control the pressure on the object.

We use following notation: The forward kinematic function that computes the position of a body part a is denoted $\phi_a^x(\mathbf{q})$; the kinematic function that computes the z -axis (normal vector) of a body part is denoted $\phi_a^z(\mathbf{q})$.

- We reuse the old idea of navigating to an object driven by the forces in a potential field. With smaller values of $f(x)$, the speed of approaching the object should be reduced, eventually becoming zero at the surface, where $f(x) = 0$. For this we introduce the hand position task variable $\mathbf{y}_h \in \mathbb{R}^3$, $\mathbf{y}_h = \phi_h^x(\mathbf{q})$.
- In order to allow the fingers to grasp an object, the wrist needs to approach the object with the front side oriented to it. In the previous section, we showed that the negative gradient points to the object. Here we require the wrist normal to align with $-\nabla f(x)$. After taking care of that in the beginning, the relevance of that argument decreases with approaching the object, i.e. with smaller ISP values, where violating it in favour of better finger configuration can be favourable. We introduce the hand axis task variable $y_p \in \mathbb{R}$, $y_p = \langle \phi_h^z(\mathbf{q}), \nabla f(\mathbf{y}_h) \rangle$ with target $y_p^* := -1$.
- For the grasp, the fingers need to touch the surface. This can be expressed as requiring that the implicit surface value at the fingers is zero. For this we introduce a task variable $\mathbf{y}_s \in \mathbb{R}^3$, $\mathbf{y}_s = (f(\phi_i^x(\mathbf{q})))_{i \in \text{tips}}$ with constant target $\mathbf{y}_s^* := (0)$. i indicates any of the 3 finger tips.

¹this is what we call *relative importance* in the beginning of section III

- When the *finger tips point to their common centroid*, we expect to encounter stable grasp (analogous to force closure). In a three-fingers setting the sum of the normals of the finger tips lives in $[0, 3]$, with zero reached when the fingers oppose perfectly. In other words, zero sum of normals is necessary condition for opposing fingers. We introduce the task variable $y_o \in \mathbb{R}$, $y_o = \left| \sum_{i \in \text{tips}} \phi_i^z(\mathbf{q}) \right|$ with constant target $y_o^* := 0$.
- *Large contact area* is achieved when the tangent of the shape is aligned with the finger tangent at the contact point, i.e., the two normals have to cancel. Since the normals of the surface coincide with the negative gradient on the surface, we insist on the fingertip's normal and the ISF value at the fingertip to be aligned. The nearest to the surface, the more the orientation of the fingers matter, thus we make this precision negatively proportional to $f(x)$. We introduce the task variable $\mathbf{y}_f \in \mathbb{R}$, $\mathbf{y}_f = (\langle \nabla f(\phi_i^x(\mathbf{q})), \phi_i^z(\mathbf{q}) \rangle)_{i \in \text{tips}}$ with constant target $\mathbf{y}_f^* := (-1)$.

In addition, we have task variables to avoid collisions and joint limits. In our controller, all velocity precisions ν_i are zero and we have no velocity targets in the task spaces.

The accumulated loss introduced by not satisfying the constraints poses an optimisation problem within the robot configuration space. The loss of a TV is given by the discrepancy between current and target value weighted by the precision, $l_a = \rho_{y_a}(y_a - y_a^*)^2$.

Algorithm 1 Grasping with ISF.

loop

$\mathbf{y}_h = \phi_h^x(\mathbf{q})$ {current position of palm center}
 $\mathbf{y}_h^* := \mathbf{y}_h - \nabla f(\mathbf{y}_h)$ {to-be position}
 $y_p = \langle \phi_p^z(\mathbf{q}), \nabla f(\mathbf{y}_h) \rangle$ {current orientation}
 $y_p^* := -1$ {to-be orientation}
 $\mathbf{c} = (0)$ {centroid of tip normals}

for k in fingers **do**

$\mathbf{y}_{k,f} = \langle \nabla f(\phi_k^x(\mathbf{q})), \phi_k^z(\mathbf{q}) \rangle$ {orientation}
 $\mathbf{y}_{k,f}^* := -1$ {target orientation}
 $\mathbf{y}_{k,s} = f(\phi_k^x(\mathbf{q}))$ {value at k -th finger}
 $\mathbf{y}_{k,s}^* := 0$ {target value}
 $\mathbf{c} = \mathbf{c} + \phi_k^z(\mathbf{q})$ {centroid of tip normals}
 set ρ_{k,y_f} and $\rho_{k,y_s} \propto (1 - f(\phi_k^x(\mathbf{q})))$

end for

$y_o = |\mathbf{c}|$
 $y_o^* := 0$ {make finger oppose}
 set $\rho_{y_h}, \rho_{y_p}, \rho_{y_o} \propto f(\phi_h^x(\mathbf{q}))$
 $\mathbf{q}' := \phi_{SOC}^{-1}(\mathbf{y})$ {satisfy constraints optimally}
 $\mathbf{q} := \mathbf{q}'$ {move}

end loop

Algorithm 1 defines the feedback control loop of the reach and grasp controller. Typically, the trajectories generated by the algorithm start with a motion towards object and simultaneously putting the wrist to point with the finger side to the object. When the fingers are near the object, the ISF falls exponentially and relative importance of finger




Fig. 5. Random object (top left); belief improvement with new observations (black dots).




Fig. 6. The median and upper and lower 5% quantile for the precision of the shape estimation as function of the number of (random) observations provided.

orientation rises with the consequence that the fingers open. Then the fingers continue to approach the object guided by the potential field.

The developed approach is not immune against local minima, which is the case with all approaches relying on local information only. Nevertheless it can find its application as a fast approximation where whole trajectory optimization is not desired or infeasible or can serve as initialization for optimisation approaches.


IV. EXPERIMENTS

A. Surface Estimation

The estimation algorithm was validated using synthetic tactile data for a random generated object: First, we generate an implicit surface representation by sampling from a GP. This yields a random object. Then, we select randomly points on the surface of that object and use them as input to a GP. The resulting mean of the GP posterior gives the estimation of the object. Refer to figure 5 for snapshots from this procedure. The first image shows the random object which is to be estimated. Starting with one observation, second image, the belief about the object improves progressively with new observations. Black dots mark observed points on the surface.

Figure 6 shows a plot of the accuracy of the estimation when using different number of observations. The quantiles for particular observation count (on the x -axis) are computed from 300 randomly generated objects of different sizes. The value on the y axis is the similarity of true and estimated object. This is obtained by measuring the shared volume of estimated and true object, subtracting the falsely estimated volume (inside the estimated, but not in the true). This absolute value is normalized by the volume of the true object.

Fig. 7. Experimental setup: Schunk 7DoF arm, Schunk 7DoF hand, Bumblebee stereo camera, coloured objects with various shapes.



B. Grasp Controller

For the evaluation of our controller we grasp real objects by using the Schunk Light Weight Arm with 7DoF and the attached Schunk Dextrous Hand with 7DoF and tactile sensors, see Figure 7. Part of the setup is also a Bumblebee stereo camera which is used to localize and provide surface points of the objects via stereo triangulation. We use a box, a cylindrical can and a ball, and place them – one at a time – on a table in front of the robot. Every object is uniformly coloured to ease the visual perception: simple hue threshold segmentation delivers a contour which fits one of the predefined shapes. Once the object is localized, we create a parametric ISP and run the grasp controller with it.

The video submitted with this document presents several grasp trajectories for objects placed at different locations on the table. The arm starts with fast movement towards the object and orients the wrist accordingly. When nearing the object, the speed reduces, the fingers open and align with the shape. This can be seen when comparing the ball and the cylinder. The fingers follow the symmetrical shape of the ball, whereas for cylinder they stay aligned with the vertical wall. Eventually, the fingers get in contact with the surface and the skin pressure rises. When it goes beyond given threshold, the controller stops and the robot lifts the object as an indication for successful grasp.

V. CONCLUSION

In this paper we present a novel object representation for robotic grasping and sensor fusion based on implicit surfaces and couple it with a robot movement control system. The overall scheme is able to approximate a large variety of object shapes, and to achieve fluent reach and grasp movements using a redundant robot with a tactile multi-finger hand. The novel contributions of this paper are

- an implicit surface description of the object’s shape which is learnt using a Gaussian process. This allows for an incremental shape approximation, consecutively incorporating sensor data from different modalities, such as visual, laser or tactile. The uncertainties of the measurements are elegantly translated into a shape uncertainty.
- a set of generic task variables that are associated with the constructed ISP and constitute to a robust reach-and-grasp movement control system. The proposed variables have strong generalization capabilities, for instance being able to grasp an object at locations it has not explored before, or to grasp objects of similar shape.

The proposed methods have been evaluated on a 14-DoF robot with a tactile 3-finger hand (7-DoF arm, 7-DoF hand) in a set of reach-grasp experiments. The results show good approximation capabilities for the object shape, and a robust and smooth movement behavior.

Future challenges include integration of the grasping and estimation tasks in a real world scenario with more relevant objects. The scenario should demonstrate the ability of GP ISF to fuse multimodal sensory information – in a form directly usable for grasping by the reach-and-grasp controller.

REFERENCES

- [1] D. Berenson, R. Diankov, K. Nishiwaki, S. Kagami, and J. Kuffner, “Grasp planning in complex scenes,” in *Proceedings of the IEEE-RAS/RSJ International Conference on Humanoid Robots*, 12 2007.
- [2] J. Blinn, “A generalization of algebraic surface drawing,” *ACM Transactions on Graphics (TOG)*, vol. 1, no. 3, pp. 235–256, 1982.
- [3] J. Bloomenthal, “Bulge elimination in implicit surface blends,” in *Implicit Surfaces*, vol. 95, 1995, pp. 7–20.
- [4] B. Curless and M. Levoy, “A volumetric method for building complex models from range images,” in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. ACM, 1996, pp. 303–312.
- [5] A. Girard, C. Rasmussen, J. Candela, and R. Murray-Smith, “Gaussian process priors with uncertain inputs-application to multiple-step ahead time series forecasting,” *Advances in Neural Information Processing Systems*, pp. 545–552, 2003.
- [6] R. Haschke, J. Steil, I. Steuwer, and H. Ritter, “Task-oriented quality measures for dextrous grasping,” in *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 6 2005, pp. 689 – 694.
- [7] O. Khatib, “Real-time obstacle avoidance for manipulators and mobile robots,” *The International Journal of Robotics Research*, vol. 5, no. 1, p. 90, 1986.
- [8] D. Meagher, “Geometric modeling using octree encoding,” *Computer Graphics and Image Processing*, vol. 19, no. 2, pp. 129–147, 1982.
- [9] A. Miller, S. Knoop, H. Christensen, and P. Allen, “Automatic grasp planning using shape primitives,” in *Proceedings of the IEEE International Conference of Robotics and Automation (ICRA)*, 2003, pp. 1824 – 1829.
- [10] Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H. Seidel, “Multi-level partition of unity implicits,” in *ACM SIGGRAPH 2005 Courses*. ACM, 2005, p. 173.
- [11] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*. Springer, 2006.
- [12] A. Saxena, J. Driemeyer, and A. Ng, “Robotic grasping of novel objects using vision,” *The International Journal of Robotics Research*, vol. 27, no. 2, p. 157, 2008.
- [13] D. Schwammkrug, J. Walter, and H. Ritter, “Rapid learning of robot grasping positions,” in *Proceedings of the International Symposium on Intelligent Robotics Systems (SIRS)*, 1999.
- [14] F. Steinke, B. Scholkopf, and V. Blanz, “Support vector machines for 3D shape processing,” in *Computer Graphics Forum*, vol. 24, no. 3, 2005, pp. 285–294.
- [15] R. Suárez, M. Roa, and J. Cornellà, “Grasp quality measures,” *Universitat Politècnica de Catalunya, Institut d’Organització i Control de Sistemes Industrials*, Tech. Rep. IOC-DT-P 2006-10, 2006.
- [16] M. Toussaint, “A Bayesian view on motor control and planning,” in *From motor to interaction learning in robots*, O. Sigaud and J. Peters, Eds. Springer, 2010.
- [17] M. Toussaint, N. Plath, T. Lang, and N. Jetchev, “Integrated motor control, planning, grasping and high-level reasoning in a blocks world using probabilistic inference,” in *Proc. of the Int. Conf. on Robotics and Automation (ICRA 2010)*, 2010.
- [18] O. Williams and A. Fitzgibbon, “Gaussian process implicit surfaces,” *Gaussian Proc. in Practice*, 2007.
- [19] H. Zha, T. Hoshida, and T. Hasegawa, “A recursive fitting-and-splitting algorithm for 3-D object modeling using superquadrics,” in *Fourteenth International Conference on Pattern Recognition, 1998. Proceedings*, vol. 1, 1998.