GP-GPIS-OPT: Grasp Planning With Shape Uncertainty Using Gaussian Process Implicit Surfaces and Sequential Convex Programming

Supplemental Material

Jeffrey Mahler[1], Sachin Patil[1], Ben Kehoe[2], Jur van den Berg[3], Matei Ciocarlie[4], Pieter Abbeel[1], Ken Goldberg[1]

This file is a supplement to the ICRA paper "GP-GPIS-OPT: Grasp Planning With Shape Uncertainty Using Gaussian Process Implicit Surfaces and Sequential Convex Programming." We refer the reader to the original paper for notation and the problem description [8]. Our dataset, references to our code, and videos of our experiments can be found at `http://rll.berkeley.edu/icra2015grasping/`.

I. GPIS Surface Normals

In this section we describe how to predict surface normals using a GPIS. Formally, the surface normal of a GPIS $f \sim \mathcal{N}(\mu, \Sigma)$ at a spatial location $\mathbf{x}$ is $\mathbf{n}(\mathbf{x}) = \frac{\bigtriangledown f(\mathbf{x})}{\|\bigtriangledown f(\mathbf{x})\|_2}$. The derivative of a Gaussian process is another Gaussian process, and therefore the gradient $\bigtriangledown f(\mathbf{x})$ can also be predicted using Gaussian process regression [10], [3]:

$$\bigtriangledown f(\mathbf{x}) = \bigtriangledown \left(m(\mathbf{x}) + k(\mathcal{X}, \mathbf{x})^\intercal (K + \sigma_m^2 I)^{-1} (\mathbf{y} - m(\mathcal{X})) \right)$$

This requires derivatives of the kernel (covariance) function. Following Dragiev et al. [3], we use the squared exponential kernel covariance function:

$$k(\mathbf{x}_i, \mathbf{x}_j) = C \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\ell^2} \right) + \sigma_m^2 \delta_{ij}$$

which specifies the correlation of the signed distance between two spatial points. This covariance depends on a scale $C \in \mathbb{R}$, bandwidth $\ell \in \mathbb{R}$, and measurement noise $\sigma_w^2 \in \mathbb{R}$, which we set using maximum-likelihood estimation [11].

The derivatives of this covariance function are:

$$k\left(\frac{\partial}{\partial \mathbf{x}} \mathbf{x}_i, \mathbf{x}_j \right) = \frac{(\mathbf{x}_j - \mathbf{x}_i)}{\ell^2} k(\mathbf{x}_i, \mathbf{x}_j)$$

$$k\left(\frac{\partial}{\partial \mathbf{x}} \mathbf{x}_i, \frac{\partial}{\partial \mathbf{x}} \mathbf{x}_j \right) = \sigma_m^2 \delta_{ij} + \frac{1}{\ell^2} k(\mathbf{x}_i, \mathbf{x}_j) + \frac{(\mathbf{x}_j - \mathbf{x}_i)(\mathbf{x}_j - \mathbf{x}_i)^\intercal}{\ell^4} k(\mathbf{x}_i, \mathbf{x}_j)$$

[1]Department of EECS; {jmahler, sachinpatil, pabbeel, goldberg}@berkeley.edu
[2]Department of ME; benk@berkeley.edu
[1-2] University of California, Berkeley, CA, USA
[3]Google Inc., CA, USA jurvandenberg@gmail.com
[4]Department of ME, Colombia University, NY, USA matei.ciocarlie@columbia.edu

We can use these definitions to form a joint covariance matrix for gradient prediction

$$\tilde{K}(\mathcal{X}_i, \mathcal{X}_j) = \begin{bmatrix} K(\mathcal{X}_i, \mathcal{X}_j) & K(\mathcal{X}_i, \bigtriangledown \mathcal{X}_j) \\ K(\bigtriangledown \mathcal{X}_i, \mathcal{X}_j) & K(\bigtriangledown \mathcal{X}_i, \bigtriangledown \mathcal{X}_j) \end{bmatrix}$$

Then, given gradient observations of an SDF $\bigtriangledown \mathbf{y}$ at spatial locations $\mathcal{X}$, we can predict both the signed distance and gradient at a point $\mathbf{x}$ as:

$$\begin{bmatrix} f(\mathbf{x}) \\ \bigtriangledown f(\mathbf{x}) \end{bmatrix} = \begin{bmatrix} m(\mathbf{x}) \\ \bigtriangledown m(\mathbf{x}) \end{bmatrix} + \tilde{K}(\mathcal{X}, \mathbf{x})^\intercal (\tilde{K}(\mathcal{X}, \mathcal{X}))^{-1} \begin{bmatrix} \mathbf{y} - m(\mathcal{X}) \\ \bigtriangledown \mathbf{y} - \bigtriangledown m(\mathcal{X}) \end{bmatrix}$$

II. Quality Approximation

Our goal is to find a grasp g that maximizes P_F given a GPIS mean and variance function μ, σ^2:

$$g^* = \max_{g \in \mathcal{G}} P\left(Q_F(\hat{g}, f) > 0 \mid g, \mu, \sigma^2 \right) . \quad (1)$$

Direct optimization of Equation (1) requires a computationally expensive integration to evaluate the objective for each grasp. However, the probability of force closure on the the mean SDF μ, $\tilde{P}_F(g, \mu)$, may be a reasonable approximation of $P_F(g, \mu, \sigma^2)$ under the following conditions:

1) The center of mass $\mathbf{z}$ is within δ of its expectation $E[\mathbf{z}]$ with high probability, $P\left(\|\mathbf{z} - E[\mathbf{z}]\|_2 > \delta\right) \approx 0$ for some $\delta > 0$.
2) The shape uncertainty along grasp approach $\mathbf{v}$ up to the contact locations $\tilde{\mathbf{c}}_i$ on the mean shape is less than some small $\tau > 0$, $\sigma^2(\tilde{\mathbf{c}}_1) < \tau$.

Consider a desired grasp g and GPIS mean μ and variance σ^2 as defined in Section V of the original paper [8]. Let $\tilde{c} = (\tilde{\mathbf{c}}_1, \tilde{\mathbf{c}}_2)$ denote the contact points for a grasp g on the mean SDF μ, determined using a search over a discrete set of points along the line of action. Let $A(f)$ denote the event that the parallel jaws contact an SDF f before reaching $\tilde{\mathbf{c}}_1$ when attempting g and passing over a discrete set of sample points. Also, let $B(f)$ denote the event that $\tilde{\mathbf{c}}_1$ is not a zero crossing of SDF sample f according to the definition of Section **??**, $|f(\tilde{\mathbf{c}}_1)| > \epsilon$.

Formally, our first assumption is that $P(A(f)) < \delta$ for some $\delta \in (0, 1)$ and that $\sigma^2(\tilde{\mathbf{c}}_1) < \tau^2$ for some $\tau \in \mathbb{R}$

and $\tau << \epsilon$. Under these assumptions, for any SDF f we have $|f(\mathbf{c}_1)| \leq |\mu(\tilde{\mathbf{c}}_1) \pm 1.96\sigma(\tilde{\mathbf{c}}_1)| < \epsilon + 2\tau \approx \epsilon$ with approximately 95% probability, since $\mathcal{N}(\mu(\tilde{\mathbf{c}}_1), \sigma^2(\tilde{\mathbf{c}}_1))$ is a 1-dimensional Gaussian [9]. Therefore $\tilde{\mathbf{c}}_1$ is still likely to be a zero crossing for other sample SDFs from the GPIS. Furthermore, the probability that we do not contact the surface at $\tilde{\mathbf{c}}_1$ when attempting grasp g is:

$$\begin{aligned} P(A(f) \cup B(f)) &\leq P(A(f)) + P(B(f)) \\ &\leq \delta + 0.05 \end{aligned}$$

Therefore $\tilde{\mathbf{c}}_1$ will be the computed contact point for approximately $(0.95 - \delta)$ of SDFs sampled from the GPIS.

Our second assumption is that $\mathbf{z} \approx E[\mathbf{z}|\mu, \sigma^2]$, which makes $\tilde{P}_F(g, \mu)$ a reasonable approximation of $P_F(g, \mu)$. Empirically we found that the variance of the center of mass is less than the grid resolution for 7 out of 8 of the test objects.

We encourage τ to be small using by penalizing the uncertainty σ^2 by $\lambda \in \mathbb{R}, \lambda > 0$.

$$\underset{g \in \mathcal{G}}{\text{maximize}} \ \tilde{P}_F(g, \mu) - \lambda(\sigma^2(\tilde{\mathbf{c}}_1) + \sigma^2(\tilde{\mathbf{c}}_2)) \qquad (2)$$

High values of λ, which correspond to smaller τ, may increase the accuracy of the approximation for selected grasps but may discard grasps with high P_F. Small values of λ may increase the set of possible grasps but the approximation will become increasingly inaccurate. In practice λ can be set using cross-validation over a set of validation shapes, similar to choosing a regularization penalty in regression models [1], [5]. The λ penalty encourages GP-GPIS-OPT to avoid regions of uncertainty, unlike past algorithms that use GPIS to encourage exploration [4].

III. Additional Experiments

This section describes additional details on experiments on GP-GPIS-OPT

A. Rate of Convergence

Fig. 1 shows the probability of force closure P_F for grasps chosen by GP-G (the GP-GPIS-OPT algorithm) and GP-P, ranking a set of grasps using Monte-Carlo integration, versus the number of random grasp samples on Object D, the toy plane. For GP-P the number of samples refers to the number of randomly selected grasps to be ranked, and for GP-G the number of samples refers to the number of random initializations. The parameters of the algorithms were an approach uncertainty of $\sigma_g = 1.5$ and friction coefficient $\gamma = 0.4$, chosen to eliminate grasps with $P_F = 0.99$. We see that GP-G plans grasps with higher P_F for the sample number of random initializations, and in only 2 iterations plans a grasp with higher P_F than the grasp planned by GP-P over all samples. The convergence here is not necessarily indicative of runtime in seconds per sample, as GP-P requires a Monte-Carlo integration over samples of shape and grasp perturbations and GP-G runs Sequential Convex Programming.

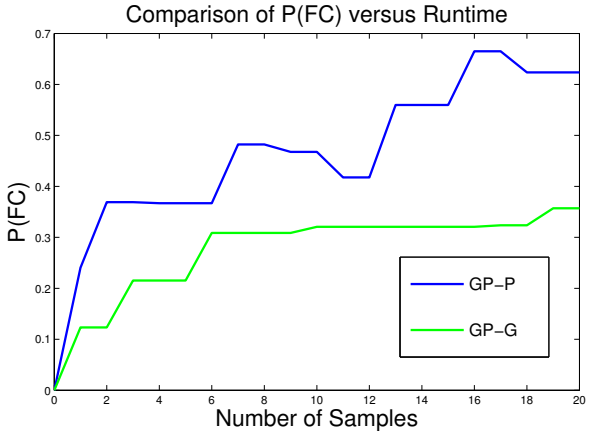

Fig. 1. Comparison of the probability of force closure P_F versus number of samples for Object D, the toy plane, averaged over 3 random sample sequences. We compare grasps planned by GP-P, which plans the grasp with the highest probability of force closure P_F from random grasp samples, and GP-G, which plans grasps using GP-GPIS-OPT. The number of samples for GP-G is the number of random initializations for the algorithm.

B. Scaling with Shape Resolution

We also studied the time complexity of our algorithm with respect to the resolution of the GPIS. We compared GP-GPIS-OPT with GP-P, the method of using Monte-Carlo integration to evaluate and rank a set of 1000 grasps [2], [6], [7]. Specifically, we compare the average runtime of our algorithm for a single initial grasp with the average time to sample 1000 shapes, a bottleneck in the Monte-Carlo evaluation of P_F. We do not attempt to compare the number of samples necessary to achieve the grasp with maximum P_F for either method because it can vary significantly based on the ordering of the random initial grasps.

The trend of the runtimes for each algorithm over increasing shape resolution for Object B, the toy plane, is plotted in Fig. 2. We see that although the method of ranking random grasps by evaluating P_F using Monte-Carlo integration performs better for low-dimensional grids, our method has a lower runtime per initialization once the grid size exceeds 55×55 and appears to have a roughly quadratic scaling with grid dimension M versus the $O(M^6)$ complexity of shape sampling. Our current implementation is also not currently optimized for speed and may experience significant speed increases when designed for high-performance. This suggests that when the space of possible grasps is larger, our algorithm may be the better choice. This includes selecting parallel jaw grasps on 3D GPIS models or optimizing over the many joint angles of a multi-fingered gripper.

C. Sensitivity to Approach Noise

We qualitatively compared grasps chosen by our algorithm with different patterns of shape uncertainty and uncertainty in the grasp approach, σ_g^2, on Object D, the metal toy plane. We compare grasps selected with $\sigma_g^2 = 0.0$ (no grasp uncertainty), $\sigma_g^2 = 1.0$, and $\sigma_g^2 = 5.0$ in Fig. 3. With smaller levels of uncertainty, the algorithm favors grasping the endpoints of the plane where there is little uncertainty

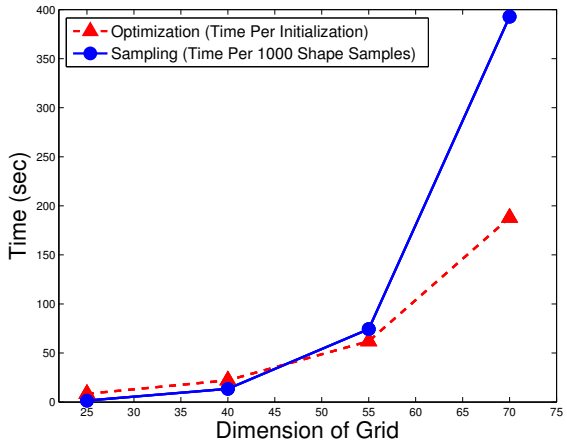

Fig. 2. Comparison of the runtime in seconds versus the dimension of the 2D GPIS grid in pixels for our algorithm versus ranking an exhaustive set of grasps using Monte-Carlo integration. The runtime for our algorithm is measured by the average optimization time per initial sample in the inner loop of our algorithm versus the time to sample 1000 shapes, a bottleneck in the Monte-Carlo evaluation of P_F. Both quantities shown here are for Object B, the toy plane. Our method exhibits better scaling with increasing shape resolution, suggesting that a speed-optimized implementation of our algorithm could be more appropriate for selecting robust grasps on 3D objects in time-sensitive applications.

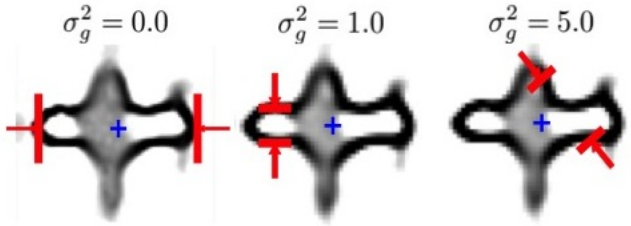

Fig. 3. Comparison of grasps chosen by our algorithm on Object B, the toy plane, with varying levels of grasp approach uncertainty σ_g^2.

but also little room for error. With increasing uncertainty the algorithm selects a grasp the plane at the front of the plane since the there is a greater surface area tangent to the grasp. At high levels of uncertainty the algorithm cannot reliably fit the parallel jaws on the front of the plane without contacting the wings, and thus selects a grasp near the wings of the plane.

REFERENCES

[1] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2009.

[2] V. N. Christopoulos and P. Schrater, "Handling shape and contact location uncertainty in grasping two-dimensional planar objects," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2007, pp. 1557–1563.

[3] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2011, pp. 2845–2850.

[4] ——, "Uncertainty aware grasping and tactile exploration," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2013, pp. 113–119.

[5] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, 1979.

[6] B. Kehoe, D. Berenson, and K. Goldberg, "Estimating part tolerance bounds based on adaptive cloud-based grasp planning with slip," in *Proc. IEEE Conf. on Automation Science and Engineering (CASE)*. IEEE, 2012, pp. 1106–1113.

[7] ——, "Toward cloud-based grasping with uncertainty in shape: Estimating lower bounds on achieving force closure with zero-slip push grasps," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 576–583.

[8] J. Mahler, S. Patil, B. Kehoe, J. van den Berg, M. Ciocarlie, P. Abbeel, and K. Goldberg, "Gp-gpis-opt: Grasp planning under shape uncertainty using gaussian process implicit surfaces and sequential convex programming," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2015.

[9] D. F. Morrison, "Multivariate statistical methods. 3," *New York, NY. Mc*, 1990.

[10] E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C. E. Rasmussen, "Derivative observations in gaussian process models of dynamic systems," 2003.

[11] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," *Gaussian Proc. in Practice*, 2007.