Investigating geometry-aware Acquisition Functions for Joint Shape Reconstructions and Grasp Improvement

Untersuchung geometriebasierter Akquisitionsfunktionen für die gemeinsame Formrekonstruktion und Greifoptimierung

Master thesis in the department of Computer Science by Vojin Vlačić
Date of submission: August 14, 2026

1. Review: Ph.D Cristiana Miranda de Farias
2. Review: MSc. Tim Schneider
3. Review: Prof. Jan Peters
Darmstadt



TECHNISCHE UNIVERSITÄT DARMSTADT



Erklärung zur Abschlussarbeit gemäß § 22 Abs. 7 APB TU Darmstadt

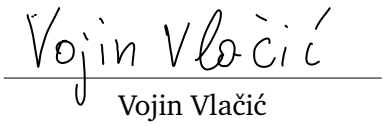
Hiermit erkläre ich, Vojin Vlačić, dass ich die vorliegende Arbeit gemäß § 22 Abs. 7 APB der TU Darmstadt selbstständig, ohne Hilfe Dritter und nur mit den angegebenen Quellen und Hilfsmitteln angefertigt habe. Ich habe mit Ausnahme der zitierten Literatur und anderer in der Arbeit genannter Quellen keine fremden Hilfsmittel benutzt. Die von mir bei der Anfertigung dieser wissenschaftlichen Arbeit wörtlich oder inhaltlich benutzte Literatur und alle anderen Quellen habe ich im Text deutlich gekennzeichnet und gesondert aufgeführt. Dies gilt auch für Quellen oder Hilfsmittel aus dem Internet.

Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Mir ist bekannt, dass im Falle eines Plagiats (§ 38 Abs. 2 APB) ein Täuschungsversuch vorliegt, der dazu führt, dass die Arbeit mit 5,0 bewertet und damit ein Prüfungsversuch verbraucht wird. Abschlussarbeiten dürfen nur einmal wiederholt werden.

Bei einer Thesis des Fachbereichs Architektur entspricht die eingereichte elektronische Fassung dem vorgestellten Modell und den vorgelegten Plänen.

Darmstadt, 14. August 2026



Vojin Vlačić

Abstract

When only part of an object's surface is observed, finding contact configurations that provide force closure is challenging due to incomplete geometric information. In the unconstrained contact-search space considered in this thesis, optimizing grasp quality alone provides no direct guidance towards the unknown object surface. To address this problem, we introduce a geometry-aware acquisition function for Bayesian optimization that combines information from a Gaussian Process Implicit Surface (GPIS) with a separate surrogate for the probability of force closure (PFC). The GPIS is updated as new geometric observations are collected, allowing surface exploration and grasp search to proceed within the same optimization process.

We evaluate the proposed method in simulation on five objects and compare it with only surrogate-guided acquisition strategies. The proposed method consistently finds high-PFC contact configurations across all tested objects. Additional experiments show that the method remains effective when hand-inspired relationships are introduced between the contacts. These results demonstrate the benefit of direct geometric guidance for multi-contact grasp search on partially observed objects.

Zusammenfassung

Wenn nur ein Teil der Oberfläche eines Objekts beobachtet wird, ist es aufgrund unvollständiger geometrischer Informationen schwierig, Kontaktkonfigurationen zu finden, die eine Kraftschließung gewährleisten. In dem in dieser Arbeit betrachteten, uneingeschränkten Kontakt-Suchraum liefert die Optimierung der Greifqualität allein keine direkten Hinweise auf die unbekannte Objektoberfläche. Um dieses Problem anzugehen, führen wir eine geometriebewusste Erfassungsfunktion für die bayessche Optimierung ein, die Informationen aus einer impliziten Gauß-Prozess-Oberfläche (GPIS) mit einem separaten Surrogat für die Wahrscheinlichkeit der Kraftschließung (PFC) kombiniert. Die GPIS wird aktualisiert, sobald neue geometrische Beobachtungen gesammelt werden, wodurch die Oberflächenerkundung und die Greifsuche im Rahmen desselben Optimierungsprozesses fortgesetzt werden können.

Wir bewerten die vorgeschlagene Methode in einer Simulation an fünf Objekten und vergleichen sie mit rein surrogatgesteuerten Erfassungsstrategien. Die vorgeschlagene Methode findet bei allen getesteten Objekten durchweg Kontaktkonfigurationen mit hohem PFC-Wert. Zusätzliche Experimente zeigen, dass die Methode auch dann effektiv bleibt, wenn handinspirierte Beziehungen zwischen den Kontakten eingeführt werden. Diese Ergebnisse belegen den Nutzen einer direkten geometrischen Führung für die Mehrpunkt-Greifersuche an teilweise beobachteten Objekten.

Contents

1. Introduction **2**
- 1.1. Problem Statement . . . 3
- 1.2. Approach and Contributions . . . 3
- 1.3. Thesis Structure . . . 4

2. Background **5**
- 2.1. Gaussian Processes . . . 5
- 2.2. Implicit Surface Representations . . . 8
- 2.3. Bayesian Optimization . . . 11
- 2.4. Force Closure and Probabilistic Grasp Evaluation . . . 18

3. Related Work **23**
- 3.1. GPIS-Based Shape Reconstruction and Tactile Exploration . . . 23
- 3.2. Grasp Planning under Shape Uncertainty . . . 25
- 3.3. Bayesian Optimization for Robotic Grasping . . . 26
- 3.4. Joint Surface Exploration and Grasp Refinement . . . 27
- 3.5. Research Gap . . . 28

4. Methodology **29**
- 4.1. Partial Point Cloud Acquisition . . . 29
- 4.2. Gaussian Process Implicit Surface . . . 31
- 4.3. Geometry-Aware Bayesian Optimization for Multi-Contact Grasp Search . . . 33

5. Experiments **42**
- 5.1. Experimental Setup . . . 42
- 5.2. Geometry-Aware Acquisition Function and Baselines . . . 48
- 5.3. Evaluation Measures . . . 50
- 5.4. Experimental Results . . . 51

6. Discussion **64**
6.1. Role of Direct Geometric Guidance 64
6.2. Comparison with Baseline Acquisition Strategies 66
6.3. Difference from the Robotic Setting of de Farias et al. 69
6.4. Effect of Surface-Based Initialization 70
6.5. Sensitivity of the Geometry-Aware Acquisition Function 71
6.6. Hand-Inspired Contact Search 72

7. Conclusion **73**
7.1. Main Findings 73
7.2. Limitations and Future Work 74

A. Detailed Experimental Results **79**
A.1. Proposed Geometry-Aware Acquisition 79
A.2. Method Based on de Farias et al. 85
A.3. Hand-Inspired Contact Configuration 86
A.4. Extended Surface-Based Warm-Up 87

1. Introduction

Grasping an object is significantly more difficult when its complete geometry is not available to the robot. A robot may initially observe only a limited portion of the object due to occlusions, sensor placement, or restricted viewpoints [1]. In such situations, grasp planning can not rely on a complete object model. Instead, the robot must learn, in some way, to handle whole objects and geometric uncertainty while searching for contact locations that can produce a stable grasp.

These two objectives are closely related. Contact points should lie on the surface of the object, but being located on the surface is not sufficient for a good grasp. The contacts must also form a configuration that can resist external disturbances and provide force closure [1]. When several contact points are optimized simultaneously, the resulting search space becomes large, while only a relatively small part of it corresponds to geometrically valid and useful grasp configurations. This difficulty becomes particularly pronounced when contact locations are searched directly in Cartesian space rather than being restricted by the kinematics of a particular robotic hand. In the formulation considered in this thesis, four three-dimensional contact locations are optimized jointly, resulting in a twelve-dimensional search space in which most candidate configurations do not correspond to useful surface contacts.

A major difficulty in multi-contact grasp search is that meaningful grasp-quality information can only be obtained once candidate contacts are sufficiently close to the object surface. Evaluating force closure for contact points that do not lie on the object surface is not physically meaningful. However, if the optimization process relies only on grasp-quality observations, it has little information about how candidate contacts should move towards the unknown surface.

This thesis proposes a geometry-aware acquisition function for Bayesian Optimization of multi-contact grasps on partially observed objects. A Gaussian Process Implicit Surface (GPIS) is used to represent the evolving object geometry, while the probability of force

closure (PFC) is used to describe grasp quality [1, 2]. The proposed approach uses information from both models to guide the search for contact configurations. As new geometric observations are obtained during the search, the GPIS model is updated, providing new information about the object's surface. In this way, surface exploration and the search for force-closure contact configurations are carried out within the same optimization process.

1.1. Problem Statement

Let

$$\mathbf{x} = \left[\mathbf{c}_1^\top, \ldots, \mathbf{c}_{N_c}^\top\right]^\top \in \mathcal{X} \subset \mathbb{R}^{3N_c} \tag{1.1}$$

denote a joint contact configuration, where N_c denotes the number of contacts and $\mathbf{c}_i \in \mathbb{R}^3$ is the position of the i-th contact. For a given object with known surface geometry $\mathcal{S}$, the multi-contact grasp-search problem can be formulated as

$$\mathbf{x}^* \in \arg\max_{\mathbf{x}\in\mathcal{X}} P_{\mathrm{FC}}(\mathbf{x}) \qquad \text{subject to} \qquad \mathbf{c}_i \in \mathcal{S}, \quad i = 1, \ldots, N_c, \tag{1.2}$$

where $P_{\mathrm{FC}}(\mathbf{x})$ denotes the probability that the contact configuration achieves force closure under the considered uncertainties.

However, in the setting considered in this thesis, the complete object surface $\mathcal{S}$ is unknown. Initially, only a partial observation of the object from one side is available, from which an initial representation of the object surface is constructed. Therefore, the constraint $\mathbf{c}_i \in \mathcal{S}$ cannot be enforced directly because large parts of the true surface are still unknown. In addition, the contact locations are optimized directly within a bounded three-dimensional search region, without imposing the kinematic constraints of a robotic hand.

The problem addressed in this thesis is therefore to find joint contact configurations that lead to a high probability of force closure while continuing to collect useful geometric information so that the incomplete object representation can be progressively refined during the search.

1.2. Approach and Contributions

The proposed method addresses this problem through a geometry-aware acquisition function that combines information from the evolving GPIS model with predicted grasp

quality. The GPIS component evaluates candidate contacts using the current surface estimate and its uncertainty, while the PFC component favors configurations predicted to provide reliable force closure.

The main contributions of this thesis are:

- a geometry-aware acquisition function that combines information from the current GPIS model with a separate PFC surrogate, allowing surface exploration and force-closure contact search to be performed within the same optimization process;
- an experimental comparison with PFC surrogate-only acquisition strategies and a method based on the combined GPIS–PFC formulation of de Farias et al. [1];
- an evaluation of the proposed method with additional hand-inspired relationships between the contact points, together with an investigation of how more informative surface-based initialization affects the combined-surrogate approach.

The proposed method is evaluated in simulation on five objects with different geometries.

1.3. Thesis Structure

The rest of this thesis is structured as follows. In Chapter 2, the theoretical background required for the proposed approach, including Gaussian processes, implicit surface modeling, Bayesian optimization, and force closure, is introduced. Chapter 3 reviews related work on object reconstruction, active exploration, and grasp optimization. The GPIS and PFC components, the acquisition function, and the optimization procedure for the proposed method are presented and described in Chapter 4. Next, in Chapter 5 the experimental setup is described and the experimental results presented. Chapter 6 discusses the main findings, comparisons with the alternative approaches, the results of the experiments, and limitations of the study. Finally, Chapter 7 summarizes the main conclusions and outlines directions for future work.

2. Background

This chapter introduces the theoretical and mathematical concepts required for the method developed in this thesis. Gaussian processes and Gaussian process regression are presented, including their predictive mean, uncertainty, and computational requirements. Implicit surface representations and Gaussian Process Implicit Surfaces are then introduced as probabilistic models of partially observed object geometry. Next, Bayesian optimization and the acquisition functions used for sequential candidate selection are described. We conclude this chapter with an overview of force closure and its probabilistic extension, which serve as criteria for evaluating candidate grasp configurations under uncertainty.

2.1. Gaussian Processes

Gaussian Processes (GPs) are a probabilistic, non-parametric method for modeling distributions over unknown functions. A GP does not assume a predefined functional form described by a fixed finite set of parameters. They account for uncertainty by considering an infinite number of possible functions rather than by depending on a predefined form with only a finite number of parameters. When observations are incorporated, a posterior distribution is obtained that not only gives us the predictions for the unknown function but also quantifies the uncertainty associated with those predictions [2, 3]. They are widely used in applications such as classification, regression, and Bayesian optimization [2, 4]. This makes them valuable in problems that require accurate predictions, but also a reliable estimate of uncertainty.

More formally, a GP can be understood as a generalization of the multivariate Gaussian distribution to functions. A GP is a group of random variables where any finite subset of them has a joint multivariate Gaussian distribution, i.e., the corresponding function values are jointly Gaussian distributed [2]. In the function-modeling context, each random

variable represents the value of an unknown function evaluated at a specific input point. Each GP is defined with its mean function and covariance function.

Let $f : \mathcal{X} \to \mathbb{R}$ denote an unknown scalar-valued function, where $\mathcal{X} \subseteq \mathbb{R}^d$ can be understood as the input domain and d is the dimension of the input space. Here, $f(\mathbf{x})$ represents the value of the unknown function at a particular input $\mathbf{x} \in \mathcal{X}$. A Gaussian process prior over f is written as [2]

$$f(\mathbf{x}) \sim \mathcal{GP}\left(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\right), \tag{2.1}$$

where $m(\mathbf{x})$ is the mean function and $k(\mathbf{x}, \mathbf{x}')$ is the covariance function, commonly referred to as the kernel. Choosing the kernel is one of the fundamentals in defining the properties and behavior of a GP.

The mean function is defined as the expected value of the function at each input location:

$$m(\mathbf{x}) = \mathbb{E}\left[f(\mathbf{x})\right]. \tag{2.2}$$

On the other hand, for two input locations, the covariance function describes statistical dependence between the corresponding function values:

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}\left[\left(f(\mathbf{x}) - m(\mathbf{x})\right)\left(f(\mathbf{x}') - m(\mathbf{x}')\right)\right]. \tag{2.3}$$

Therefore, the covariance function determines which functions are considered likely under the GP prior. For example, a kernel that assigns high covariance to nearby input points expresses the assumption that the corresponding function values should be similar. On the other hand, a low covariance indicates a weaker relationship between the function values.

In most practical applications, the prior mean is assumed to be zero $m(\mathbf{x}) = 0$. The main reason for such an assumption is often adopted for mathematical convenience. However, a zero prior mean does not imply that the posterior mean must also be zero. After GP is conditioned on the observed data, the posterior mean is determined by the observations and is generally non-zero. The observed data modify the posterior distribution and shift the posterior mean toward values supported by the measurements. A non-zero mean function may also be used when prior knowledge about the general behavior of the unknown function is available.

2.1.1. Gaussian Process Regression

As introduced in Section 2.1, a Gaussian process defines a prior distribution over functions. Gaussian process regression (GPR) can be defined as the use of this prior distribution to predict an unknown continuous function from a finite set of observed input-output pairs. More formally, GPR uses a Gaussian process prior to infer an unknown continuous function from a finite set of observations. In addition to predicting the value of the function at an unobserved input, GPR provides an estimate of the uncertainty associated with that prediction [2, 3, 5].

Consider a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{n}$, where $\mathbf{x}_i \in \mathbb{R}^d$ is an input point and y_i is the output i.e. corresponding observed value. All observations are considered as noisy measurements of an underlying latent function: $y_i = f(\mathbf{x}_i) + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma_n^2)$, represents zero-mean Gaussian observation noise and σ_n^2 denotes the observation-noise variance.

Let $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_n]$ be the set of training inputs, and let $\mathbf{y} = \begin{bmatrix} y_1 & \cdots & y_n \end{bmatrix}^T$ contain the corresponding observations. We write the covariance matrix of the training inputs as $K \in \mathbb{R}^{n \times n}$, with elements $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. For an unseen test input $\mathbf{x}_*$, the covariance between the test input and the training inputs is denoted by

$$\mathbf{k}_* = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_*) & \cdots & k(\mathbf{x}_n, \mathbf{x}_*) \end{bmatrix}^T. \tag{2.4}$$

Conditioning the GP prior on the observed data yields a Gaussian posterior predictive distribution $p(f_* \mid \mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}\left(\mu(\mathbf{x}_*), \sigma^2(\mathbf{x}_*)\right)$, where $f_* = f(\mathbf{x}_*)$. Following [2, Eqs. (2.25, 2.26)] and assuming a zero prior mean, the posterior mean is given by

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^T \left(K + \sigma_n^2 I\right)^{-1} \mathbf{y}, \tag{2.5}$$

where I denotes the $n \times n$ identity matrix, and the posterior variance is

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T \left(K + \sigma_n^2 I\right)^{-1} \mathbf{k}_*. \tag{2.6}$$

The posterior mean represents the predicted value of the latent function at $\mathbf{x}_*$, while the posterior variance quantifies the uncertainty associated with this prediction. Generally, the variance is lower in regions supported by nearby observations, i.e., where we have data, and higher in regions that have not been sufficiently explored. This distinction between predicted value and predictive uncertainty is particularly important in Bayesian

optimization [4], which is introduced in a later section. Many acquisition functions, including Expected Improvement [6] and Upper Confidence Bound [7], use these quantities to determine the next point to evaluate.

2.1.2. Computational Considerations

Solving linear systems involving the covariance matrix is required to use exact GP inference. Although the predictive equations are commonly written using the inverse of the covariance matrix K_y^{-1}, it should generally not be computed explicitly in a numerical implementation. Instead of explicitly computing $(K + \sigma_n^2 I)^{-1}$, as a solution, a Cholesky decomposition is commonly applied $K_y = LL^T$, where L is a lower-triangular matrix and where $K_y = K + \sigma_n^2 I$. The required linear systems can be solved efficiently using forward and backward substitution. For an n-training-dataset, the computational complexity of GP training is approximately $\mathcal{O}\left(n^3\right)$, primarily due to the factorization of the covariance matrix. The memory requirement is approximately $\mathcal{O}\left(n^2\right)$ [2].

These requirements can become restrictive for large datasets. To improve scalability, several scalable GP methods have been developed, including sparse Gaussian processes, inducing-point methods, local Gaussian processes, and structured kernel approximations [8]. For datasets of small to moderate size, exact GPR remains attractive because it provides an analytically tractable posterior distribution and an explicit estimate of predictive uncertainty.

2.2. Implicit Surface Representations

For a three-dimensional object, we can represent its geometry by a polygonal mesh which consists of vertices, edges, and faces. This type of the representation provides direct access to the object and its surface. However, there is no information about uncertainty i.e., a conventional polygonal mesh does not inherently encode uncertainty in regions where the number of measurements is small or unavailable. Another way to represent an object implicitly is to use a continuous scalar function defined over the surrounding three-dimensional space [9, 10]. When we have partially observed objects, implicit surface representations are particularly suitable because they provide a continuous description of the surface and do not require an explicit parametrization of its topology. In addition, when combined with Gaussian Processes, they can quantify the uncertainty associated with the reconstructed shape of the object.

2.2.1. Implicit Surface

Let $\phi : \mathbb{R}^3 \to \mathbb{R}$ denote a scalar-valued implicit function. The corresponding implicit surface is defined by the following function

$$\mathcal{S} = \left\{ \mathbf{x} \in \mathbb{R}^3 \mid \phi(\mathbf{x}) = 0 \right\}. \tag{2.7}$$

Points satisfying $\phi(\mathbf{x}) = 0$, we say that they lie on the object surface. The sign of the function can additionally be used to distinguish the object interior from the surrounding free space. A widely used and accepted conventions [1, 10, 11] is

$$\phi(\mathbf{x}) = \begin{cases} > 0, & \text{if } \mathbf{x} \text{ is outside the object,} \\ = 0, & \text{if } \mathbf{x} \text{ is on the surface,} \\ < 0, & \text{if } \mathbf{x} \text{ is inside the object.} \end{cases} \tag{2.8}$$

2.2.2. Gaussian Process Implicit Surfaces

A Gaussian Process Implicit Surface (GPIS) represents an implicit surface probabilistically by placing a Gaussian process prior over the implicit function $\phi(\mathbf{x})$ [9, 10]. The implicit function is therefore modeled as

$$\phi(\mathbf{x}) \sim \mathcal{GP}\left(m_\phi(\mathbf{x}), k_\phi(\mathbf{x}, \mathbf{x}')\right), \tag{2.9}$$

where $m_\phi(\mathbf{x})$ is the prior mean function and $k_\phi(\mathbf{x}, \mathbf{x}')$ is the covariance function. The covariance function characterizes the spatial relationships among implicit function values and determines properties including the smoothness and extrapolation behavior of the reconstructed surface.

When we consider a set of implicit-surface observations $\mathcal{D}_{\text{IS}} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $\mathbf{x}_i \in \mathbb{R}^3$ is a spatial location in terms of position in space and y_i is the corresponding implicit-function value. Points that are observed on the surface are assigned $y_i = 0$. For the other points, for which we know that they lie inside or outside, they may be assigned negative or positive values such as -1 and $+1$, respectively. With this set of signed observations, we allow the model to distinguish the object interior from the surrounding free space, i.e., on the surface, outside of the surface, or inside the surface.

In terms of robotics and tactile exploration, a detected contact at $\mathbf{c}_i \in \mathbb{R}^3$ provides the surface constraint $\phi(\mathbf{c}_i) = 0$. With such a set of surface observations, we know only where the surface of the object is, but there is no indication which side of the surface corresponds to the object interior. Therefore, the zero-valued surface observations can be combined with positive and negative observations to provide information about the sides of the objects and about the sign of the implicit function. Each new observation is incorporated incrementally into the GPIS model; therefore, we update both the estimated surface and its associated uncertainty, as well as which side of the space is inside and which is outside the object [1, 10].

Following the Gaussian process regression formulation introduced in Section 2.1.1, and conditioning the implicit function on the available surface observations yields the posterior predictive distribution [1, 2]

$$p\left(\phi_* \mid \mathbf{x}_*, \mathcal{D}_{\mathrm{IS}}\right) = \mathcal{N}\left(\mu_\phi(\mathbf{x}_*), \sigma_\phi^2(\mathbf{x}_*)\right), \tag{2.10}$$

where $\phi_* = \phi(\mathbf{x}_*)$. The posterior mean and posterior variance are obtained using the Gaussian process regression equations introduced in Section 2.1.1.

The reconstructed object surface is defined as the zero level set of the posterior mean $\widehat{\mathcal{S}} = \left\{\mathbf{x} \in \mathbb{R}^3 \mid \mu_\phi(\mathbf{x}) = 0\right\}$. With the posterior mean, we get the current estimate of the object geometry. In contrast, the posterior variance $\sigma_\phi^2(\mathbf{x})$ quantifies the uncertainty associated with the estimated implicit function, i.e., how uncertain the model is about the implicit-function value at a given point and, consequently, how uncertain the model is about whether that point lies outside, on, or inside the object. In the case of the regions that are supported by nearby observations, they generally have a lower posterior variance, whereas regions containing few or no observations generally have a higher variance [1, 11]. The posterior variance of the GPIS in this case is not the geometric distance from some point to the object surface; it represents uncertainty about that estimate. The posterior mean represents the estimated implicit field. This distinction enables identification of regions where additional observations may improve the reconstructed surface.

An important property of GPIS models is that they can be updated incrementally as soon as new observations become available. For each additional observation, we recondition the Gaussian process, and with more data, we update both the posterior mean and the posterior variance. Therefore, additional information gathered during exploration can refine the estimated surface and reduce uncertainty in regions supported by new observations [1, 10].

2.3. Bayesian Optimization

Bayesian optimization (BO) [4] is a sequential optimization framework designed for objective functions that are too expensive to evaluate, lack a known analytical form, or lack gradient information. Such functions are commonly referred to as black-box objective functions. Instead of evaluating the objective function at a large number of points, which is expensive and time-consuming, Bayesian optimization constructs a probabilistic surrogate model and uses it to determine which input should be evaluated next.

The surrogate model provides both a prediction of the objective function and an estimate of the uncertainty associated with that prediction. An acquisition function that aims to propose the next point uses these quantities to balance the exploitation of currently promising regions and the exploration of regions where the surrogate model remains uncertain. After the next point has been selected by maximizing the acquisition function, the point is evaluated, the resulting observation is added to the available data, and the surrogate model is updated. This process continues until an evaluation budget is reached or another stopping criterion is met.

2.3.1. Problem Formulation and Surrogate Model

Consider an unknown objective function $f : \mathcal{X} \to \mathbb{R}$, where $\mathcal{X}$ denotes the search domain. Using a maximization convention, the optimization problem is defined as in [4]

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}). \tag{2.11}$$

On the other hand, a minimization problem can be represented in the same form by negating its objective function. More precisely, minimizing a cost function $c(\mathbf{x})$ is exactly equivalent to maximizing the reward function, which we can write as $f(\mathbf{x}) = -c(\mathbf{x})$.

At iteration t, the available observations are represented by the dataset $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^{t}$, where $\mathbf{x}_i \in \mathcal{X}$ is an evaluated input and y_i is the output, i.e. corresponding observed objective value. In the output, we also assume some part of the observation noise $y_i = f(\mathbf{x}_i) + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$.

A Gaussian process, described in detail in Section 2.1, is commonly used as the surrogate model. The main reason is that it provides both a prediction of the unknown objective

function and an estimate of the uncertainty associated with that prediction [4]. Following the Gaussian process regression formulation introduced in Section 2.1.1, the predictive distribution at a candidate input $\mathbf{x}$ is

$$f(\mathbf{x}) \mid \mathcal{D}_t \sim \mathcal{N}\left(\mu_t(\mathbf{x}), \sigma_t^2(\mathbf{x})\right), \tag{2.12}$$

where $\mu_t(\mathbf{x})$ is the posterior mean and $\sigma_t^2(\mathbf{x})$ is the posterior variance after conditioning the Gaussian process on the first t observations. The posterior mean represents the current prediction of the objective value at $\mathbf{x}$, and the posterior variance quantifies the uncertainty associated with that prediction. Bayesian optimization uses the acquisition function to decide which input to evaluate next, combining both quantities.

2.3.2. Acquisition Functions

The acquisition function assigns a score to each candidate point based on the expected utility of evaluating the objective function at that point. It is computed from the posterior distribution of the surrogate model and guides the selection of the next candidate [4, 12].

One of the earliest works on Bayesian sequential optimization can be traced to Kushner [13]. This work proposed selecting the next evaluation point based on the probability of improving upon the best solution observed so far. This method is an early example of an acquisition function.

A central aspect of acquisition functions is the balance between exploitation and exploration. The more exploitation we have, the more candidate points with promising posterior mean values are favored, whereas exploration favors points with high predictive uncertainty. Too much exploitation may cause the optimization procedure to remain permanently concentrated around a suboptimal region. In contrast, excessive exploration may spend evaluations in uncertain regions that are unlikely to contain high-quality solutions. Balancing these two competing objectives is commonly referred to as the exploration–exploitation trade-off [1, 4].

The next evaluation point is selected by maximizing the acquisition function [4, Algorithm 1]:

$$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \mathcal{X}}{\arg\max}\ \alpha_t(\mathbf{x}). \tag{2.13}$$

where $\alpha_t(\mathbf{x})$ denotes the acquisition function constructed from the surrogate-model posterior at iteration t. After selecting $\mathbf{x}_{t+1}$, the objective function is evaluated, and the resulting observation is added to the dataset $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{x}_{t+1}, y_{t+1})\}$.

Since the surrogate model is conditioned on the updated dataset, it produces an updated posterior mean and posterior variance. After the update, the acquisition function is recomputed, and the procedure is repeated. Expected Improvement and Upper Confidence Bound are two widely used acquisition functions and are described in the following subsections [6, 7].

2.3.3. Expected Improvement

Expected Improvement (EI) is an acquisition function that seeks out candidate points with the greatest promise, measuring how much their objective values could surpass the best result found so far. Expected Improvement was first proposed by Močkus and was later popularized through the Efficient Global Optimization method introduced by Jones et al. [6,14]. Figure 2.1 provides a one-dimensional illustration of the Expected Improvement acquisition function and the selected next evaluation point.

For a maximization problem, under the classical noise-free EI formulation,
let $f_t^+ = \max_{1 \le i \le t} y_i$ denote the best objective value observed during the first t evaluations. The improvement obtained by evaluating a candidate point $\mathbf{x}$ is defined as in [15]

$$I_t(\mathbf{x}) = \max\left(f(\mathbf{x}) - f_t^+ - \xi, 0\right), \tag{2.14}$$

where $\xi \ge 0$ is an optional improvement parameter and controls the trade-off between exploration and exploitation. When $\xi = 0$, any value that is greater than the current best value is understood as an improvement. Increasing ξ raises the required improvement threshold and generally encourages more exploration. On the other hand, smaller values of ξ place greater emphasis on exploitation.

Since $f(\mathbf{x})$ is actually unknown before the candidate point is evaluated, we can consider $I_t(\mathbf{x})$ to be a random variable under the Gaussian process posterior. The expected improvement is therefore defined as [4, Eq. (7)]

$$\alpha_{\mathrm{EI},t}(\mathbf{x}) = \mathbb{E}\left[I_t(\mathbf{x}) \mid \mathcal{D}_t\right]. \tag{2.15}$$

Using the Gaussian predictive distribution introduced in Section 2.3.1, assuming $\sigma_t(\mathbf{x}) > 0$, and following the works [4, 15] the standardized improvement is defined as

$$z_t(\mathbf{x}) = \frac{\mu_t(\mathbf{x}) - f_t^+ - \xi}{\sigma_t(\mathbf{x})}. \tag{2.16}$$

Therefore, the expected improvement can then be written and evaluated in closed form as

$$\alpha_{\mathrm{EI},t}(\mathbf{x}) = \left(\mu_t(\mathbf{x}) - f_t^+ - \xi\right) \Phi\left(z_t(\mathbf{x})\right) + \sigma_t(\mathbf{x}) \varphi\left(z_t(\mathbf{x})\right), \tag{2.17}$$

where $\Phi(\cdot)$ and $\varphi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively. For $\sigma_t(\mathbf{x}) = 0$, the Expected Improvement value is defined as zero.

The first term in Equation 2.17 favors candidate points for which it is expected to exceed the best objective value observed so far. The second term accounts for predictive uncertainty and can assign a positive acquisition value to uncertain candidate points even when their posterior mean does not exceed the current best value. Although these terms are often associated with exploitation and exploration, respectively, both depend on the standardized improvement $z_t(\mathbf{x})$. Expected Improvement therefore accounts for both the probability and the possible magnitude of an improvement.

The next candidate point is selected by maximizing Expected Improvement [4, Eq. (9)]:

$$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \mathcal{X}}{\arg\max}\ \alpha_{\mathrm{EI},t}(\mathbf{x}). \tag{2.18}$$

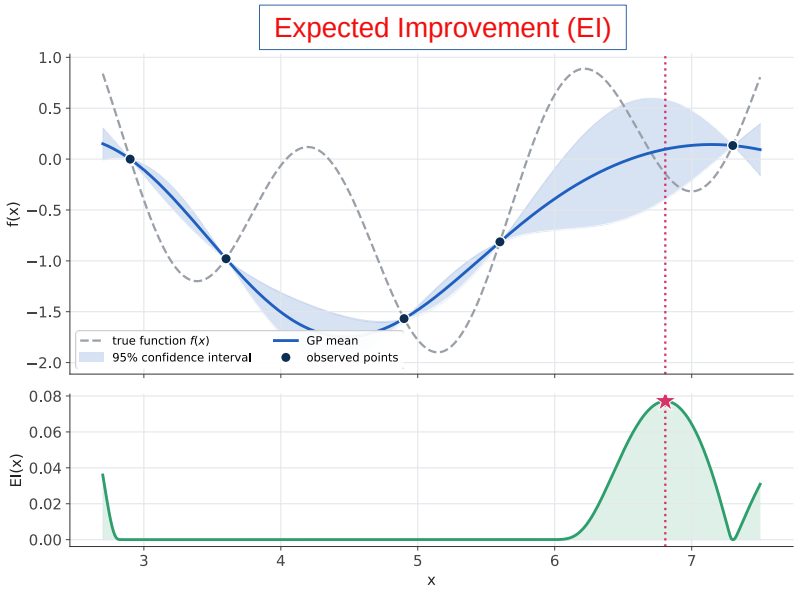


Figure 2.1.: Illustration of the Expected Improvement acquisition function in a one-dimensional Bayesian optimization example.

2.3.4. Upper Confidence Bound

Upper-confidence-bound methods originate from the multi-armed bandit literature, where decisions are made using optimistic estimates of unknown rewards [16, 17]. The Gaussian Process Upper Confidence Bound (GP-UCB) method applies this principle to the Gaussian process posterior by combining its posterior mean with its posterior standard deviation [7, Sec. 3, Eq. (6)].

Using the notation adopted in this thesis, the GP-UCB acquisition function for a maximization problem is defined as

$$\alpha_{\mathrm{UCB},t}(\mathbf{x}) = \mu_t(\mathbf{x}) + \kappa_t \sigma_t(\mathbf{x}), \tag{2.19}$$

where $\mu_t(\mathbf{x})$ and $\sigma_t(\mathbf{x})$ denote the posterior mean and posterior standard deviation, respectively, after conditioning the Gaussian process on the observations available at iteration t. The parameter $\kappa_t > 0$ controls the balance between exploration and exploitation.

The original GP-UCB formulation is commonly written using $\sqrt{\beta_t}$ instead of κ_t, such that $\kappa_t = \sqrt{\beta_t}$ [7]. The two parameterizations therefore represent the same acquisition criterion.

The posterior mean term $\mu_t(\mathbf{x})$ favors candidate points with high predicted objective values and therefore enforces more exploitation. The uncertainty term $\kappa_t \sigma_t(\mathbf{x})$ increases the acquisition value of candidate points in regions where the surrogate model is uncertain, and with an increased uncertainty term, we promote more exploration. A larger value of κ_t means we place greater relative importance on predictive uncertainty and generally encourage exploration. With a smaller κ_t value, we give greater relative importance to the posterior mean and therefore encourage exploitation. The parameter may be fixed throughout the optimization procedure or changed between iterations. In the theoretical GP-UCB formulation, the corresponding parameter β_t is commonly selected as an iteration-dependent confidence parameter [7].

Using a maximization convention, the next candidate point is selected by maximizing the GP-UCB acquisition function:

$$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \mathcal{X}}{\arg\max}\ \alpha_{\mathrm{UCB},t}(\mathbf{x}). \tag{2.20}$$

The GP-UCB value represents an optimistic estimate of the unknown objective function. A candidate point can obtain a high acquisition value either because its posterior mean is high or because its predictive uncertainty is large. This allows us to evaluate promising regions and investigate those that remain insufficiently explored. Figure 2.2 illustrates how the posterior mean and predictive uncertainty determine the Upper Confidence Bound acquisition value in a one-dimensional example.

The cost function itself can be negated, and therefore the reward can be optimized using UCB. For a cost function $c(\mathbf{x})$ that is to be minimized, the corresponding Lower Confidence Bound (LCB) may be defined as

$$\mathrm{LCB}_t(\mathbf{x}) = \mu_t(\mathbf{x}) - \kappa_t \sigma_t(\mathbf{x}), \tag{2.21}$$

and the next point is selected according to

$$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \mathcal{X}}{\arg\min}\ \mathrm{LCB}_t(\mathbf{x}). \tag{2.22}$$

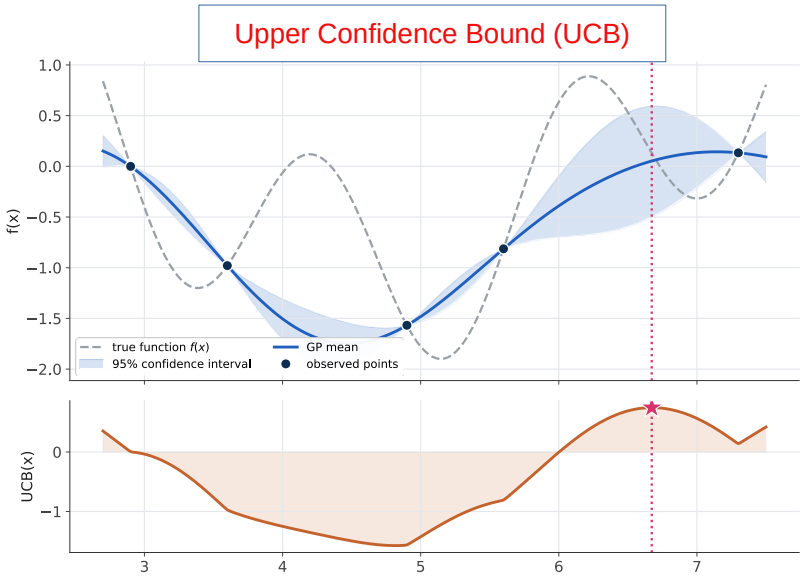


Figure 2.2.: Illustration of the Upper Confidence Bound acquisition function in a one-dimensional Bayesian optimization example.

Expected Improvement and GP-UCB both rely on the same posterior mean and predictive uncertainty, but combine them differently. While Expected Improvement measures the expected improvement over the current best observation, UCB forms an optimistic score from the posterior mean and standard deviation, sidestepping direct dependence on the current best observation.

In grasp-search applications, the input to Bayesian optimization may represent a candidate contact configuration, while the objective function represents the quality of that configuration. Force closure and its probabilistic extension, the Probability of Force Closure (PFC), provide criteria for evaluating such configurations [1]. These concepts are introduced in the following section.

2.4. Force Closure and Probabilistic Grasp Evaluation

For a contact configuration, its ability to maintain a stable grasp under external disturbances must be evaluated, especially in the presence of external forces. Force closure is a standard way to check if the allowed contact forces can balance out the external forces and moments on an object. Force closure relies on contact locations, surface normals, and friction properties. When these factors are uncertain, a probabilistic approach is appropriate for evaluating grasp stability [1, 18, 19].

2.4.1. Force Closure

Consider a contact configuration defined by the set $\mathcal{C} = \{\mathbf{c}_i\}_{i=1}^{N_c}$, where $\mathbf{c}_i \in \mathbb{R}^3$ denotes the position of the i-th contact point and N_c is the number of contacts. Let $\mathbf{n}_i \in \mathbb{R}^3$ denote the unit surface normal associated with the i-th contact, directed towards the object interior. Under the point-contact-with-friction model, although it can transmit the force, it can not directly transmit a moment. The contact force $\mathbf{f}_i$ can be decomposed into normal and tangential components [18, Ch. 5]:

$$\mathbf{f}_i = f_{n,i}\mathbf{n}_i + \mathbf{f}_{t,i}, \qquad \mathbf{n}_i^T \mathbf{f}_{t,i} = 0, \tag{2.23}$$

where $f_{n,i} \geq 0$ is the magnitude of the normal component and $\mathbf{f}_{t,i}$ is the tangential component. The Coulomb friction model requires that the tangential force must satisfy the following condition

$$\|\mathbf{f}_{t,i}\|_2 \leq \mu_i f_{n,i}, \tag{2.24}$$

where $\mu_i \geq 0$ is the coefficient of friction at the contact [18, Ch. 5]. The set of forces satisfying Equation 2.24 forms a friction cone centered around the contact normal. Forces within this cone satisfy the adopted no-slip friction constraint. When the friction cone is approximated by N_e discrete edge directions, let $\mathbf{f}_{i,j}$ denote the j-th primitive contact force at contact i, where $j = 1, \ldots, N_e$. The admissible contact forces of the discretized friction cone are given by the non-negative span of these primitive edge forces:

$$\mathcal{F}_i = \left\{ \sum_{j=1}^{N_e} \alpha_{i,j}\mathbf{f}_{i,j} \;\middle|\; \alpha_{i,j} \geq 0 \right\}. \tag{2.25}$$

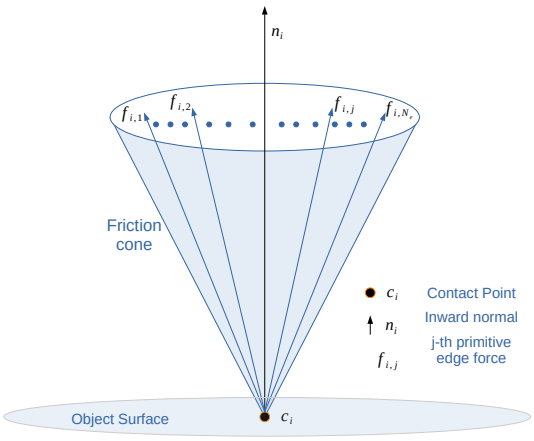


Figure 2.3.: Discretized Coulomb friction cone at contact i. The cone is approximated by N_e primitive edge forces $\mathbf{f}_{i,j}$, whose non-negative span represents the admissible contact forces.

Figure 2.3 illustrates the Coulomb friction cone and its discretization into primitive edge forces at a single contact. A contact force produces both a force and a moment on the object. Therefore, the corresponding contact wrench, expressed relative to a reference point $\mathbf{p}_{\text{ref}}$, is given by [18]

$$\mathbf{w}_{i,j} = \begin{bmatrix} \mathbf{f}_{i,j} \\ (\mathbf{c}_i - \mathbf{p}_{\text{ref}}) \times \mathbf{f}_{i,j} \end{bmatrix} \in \mathbb{R}^6. \tag{2.26}$$

The first three components of $\mathbf{w}_{i,j}$ represent the force, whereas the remaining three components represent the moment. The reference point is commonly selected as the estimated center of mass of the object. When comparing wrench magnitudes, normalization is necessary because force and moment components use different physical units [19].

The primitive wrenches obtained from the non-negative span of the edge directions of the discretized friction cones determine the wrenches available at each contact. Combining these primitive wrenches over all contacts yields the set of wrenches that the grasp can generate, referred to as the Grasp Wrench Space (GWS). Under the normalized total

contact-force constraint considered here, the GWS can be represented by the convex hull of the primitive contact wrenches [19]:

$$\mathcal{W}_{\mathrm{GWS}} = \mathrm{conv}\left(\bigcup_{i=1}^{N_c} \{\mathbf{w}_{i,j}\}_{j=1}^{N_e}\right). \tag{2.27}$$

The coefficient of friction determines the opening of each friction cone and therefore the set of admissible contact forces. Consequently, uncertainty in the friction coefficient changes the primitive contact forces and the resulting grasp wrench space.

To achieve force closure, a contact configuration must balance an arbitrary external wrench using the admissible contact forces. Under the GWS representation, this condition is satisfied when the origin lies in the interior of the grasp wrench space [18, 19]:

$$\mathbf{0} \in \mathrm{int}\left(\mathcal{W}_{\mathrm{GWS}}\right). \tag{2.28}$$

If the origin lies in the interior of the GWS, the grasp can generate wrenches in every direction around the origin and can therefore balance any sufficiently small external disturbance. Force closure provides a binary criterion for classifying a contact configuration as force-closure or non-force-closure. Ferrari and Canny introduced a quantitative grasp-quality measure based on the distance from the origin to the boundary of the GWS [19]. For a force-closure grasp, the corresponding ϵ-metric can be expressed as

$$\epsilon_{\mathrm{GWS}} = \max\left\{r \geq 0 \mid \mathcal{B}_r(\mathbf{0}) \subseteq \mathcal{W}_{\mathrm{GWS}}\right\}, \tag{2.29}$$

where

$$\mathcal{B}_r(\mathbf{0}) = \left\{\mathbf{w} \in \mathbb{R}^6 \;\middle|\; \|\mathbf{w}\|_2 \leq r\right\} \tag{2.30}$$

is a ball of radius r centered at the origin of the wrench space. More formally, ϵ_{GWS} is the radius of the largest origin-centered ball contained in the GWS.

If the value of ϵ_{GWS} is positive, it means the origin is inside the GWS and that the contact configuration satisfies force closure. A higher value shows that the system can handle a larger worst-case disturbance wrench given the force constraints, wrench normalization, and wrench-space norm [19]. This type of evaluation assumes that we know the contact positions, surface normals, friction coefficients, and reference point exactly. However,

when objects are only partially observed, these values can be uncertain. This uncertainty leads to the need for a probabilistic approach to force closure.

2.4.2. Probabilistic Force Closure

Let $\boldsymbol{\theta}$ denote the collection of quantities used to construct the grasp wrench space:

$$\boldsymbol{\theta} = \{\mathbf{c}_1, \ldots, \mathbf{c}_{N_c}, \mathbf{n}_1, \ldots, \mathbf{n}_{N_c}, \mu_1, \ldots, \mu_{N_c}, \mathbf{p}_{\text{ref}}\}. \tag{2.31}$$

Here, $\mathbf{c}_i \in \mathbb{R}^3$ denotes the position of the i-th contact point, $\mathbf{n}_i \in \mathbb{R}^3$ is the corresponding unit surface normal, μ_i is the friction coefficient at that contact, and $\mathbf{p}_{\text{ref}} \in \mathbb{R}^3$ denotes the reference point used for computing contact moments, typically the estimated center of mass of the object. The symbol N_c denotes the total number of contacts.

Having only partial object information, i.e., the object geometry and its physical properties are not known exactly, these quantities can be treated as random variables. For example, uncertainty may arise in the estimated contact positions, surface normals, friction coefficients, and center of mass [1, Sec. III-D].

The uncertainty in these parameters induces uncertainty in the friction cones, primitive contact wrenches, grasp wrench space, and the resulting value of ϵ_{GWS}. Probability of Force Closure (PFC) describes the probability that a contact configuration satisfies a specified force-closure condition. Following the probabilistic formulation used by De Farias et al., it is defined as [1, Sec. III-D]

$$P_{\text{FC}} = \Pr\left(\epsilon_{\text{GWS}}\left(\boldsymbol{\theta}\right) > \delta\right), \tag{2.32}$$

where $\delta \geq 0$ is a selected grasp-quality threshold. When $\delta = 0$, Equation 2.32 gives the probability that the contact configuration satisfies the basic force-closure condition.

The probability in Equation 2.32 can be approximated using Monte Carlo sampling. Let

$$\boldsymbol{\theta}^{(k)} \sim p\left(\boldsymbol{\theta}\right), \qquad k = 1, \ldots, N_s, \tag{2.33}$$

denote N_s independent samples of the uncertain parameters. For each sample, a new set of friction cones and contact wrenches is generated, the corresponding GWS is constructed, and the force-closure condition is evaluated. The Monte Carlo estimate of PFC is then

$$\widehat{P}_{\mathrm{FC}} = \frac{1}{N_s} \sum_{k=1}^{N_s} \mathbf{1} \left[\epsilon_{\mathrm{GWS}} \left(\boldsymbol{\theta}^{(k)} \right) > \delta \right], \tag{2.34}$$

where $\mathbf{1}[\cdot]$ is the indicator function, equal to one when the condition is satisfied and zero otherwise. This sampling-based procedure follows the general approach described by de Farias et al., where uncertain contact positions, normals, friction, and center of mass are sampled before constructing the GWS and evaluating force closure [1, Sec. III-D and Algorithm 1]. The resulting value satisfies

$$0 \leq \widehat{P}_{\mathrm{FC}} \leq 1. \tag{2.35}$$

A value close to one indicates that the contact configuration satisfies the selected force-closure condition for most sampled realizations of the uncertain parameters. On the other hand, when the value is close to zero, it indicates that force closure is rarely achieved under the assumed uncertainty model. In contrast to the deterministic force-closure criterion, PFC therefore provides a probabilistic grasp-quality score that accounts for uncertainty and can be used to compare candidate contact configurations.

3. Related Work

The problems that are considered in this work lie at the intersection of probabilistic surface reconstruction, active tactile exploration, Bayesian optimization, and the way to ensure force closure. Previous research addressed each of these topics from several different perspectives. Some of them use GP or Gaussian Process Implicit Surfaces (GPIS) to represent an object and its unknown geometry, while others additionally use the model to determine where it has uncertainty and where the surface measurement should be collected [10, 20, 21, 22]. Other approaches use probabilistic grasp-quality measures to identify how stable the grasp is under uncertainty and outside influence [1, 11]. A further line of research applies Bayesian optimization to reduce the number of expensive physical or simulated grasp evaluations required to identify robust grasp configurations [23, 24, 25].

These research areas are mostly closely related; however, they generally differ in the role assigned to geometric uncertainty. For example, in surface exploration, geometric uncertainty is commonly treated as the main quantity to be reduced. On the other hand, for grasp optimization, it is often assumed that geometry is given, and here the objective is to represent grasp quality and its robustness [11, 23, 24]. Methods that jointly refine the object model and search for stable contact configurations are less common, with the method of de Farias et al. representing a notable example [1]. The following sections review these research directions and identify the relationships among them and with the present work

3.1. GPIS-Based Shape Reconstruction and Tactile Exploration

As we already discussed in Section 2.2.2, Gaussian Process Implicit Surfaces (GPIS) representation provides a probabilistic model of object geometry, where the posterior mean defines an implicit surface estimate, and the posterior variance defines how uncertain the model is in the underlying implicit function [10, 22]. For the partially observed

objects, this combination makes GPIS models suitable because the surface estimate and its uncertainty about points and the whereabouts of the object can both be updated when new evaluations of the data (visual or tactile) measurements become available.

Dragiev et al. applied the use of GPIS models for robotic shape estimation and grasping [10]. In their work, they conditioned a Gaussian process on sensory observations, so that the zero level set of the posterior mean represents the estimated object surface. For reach-and-grasp control, they used the learned implicit shape potential and its gradient field to define task-space objectives for the robot hand. This work established an important connection between probabilistic implicit surface representations and robotic grasping. However, the posterior variance was not used to optimize the target grasp or to construct an acquisition criterion for selecting joint contact configurations.

Another work investigated how the GP uncertainty can be used to select informative tactile measurements. The method of active tactile exploration for selecting new touch using the uncertainty of a probabilistic surface model was proposed by Yi et al. [20]. Their implementation used an explicit GP surface representation rather than a GPIS model [20]. Rather than using a fixed grid to sample the object surface, their approach directed the robot to take tactile measurements in areas where the surface estimate is uncertain. The method was evaluated using a real robot, and the active-touch strategy reconstructed object surfaces more efficiently than random tactile sampling. The main objective was to reduce uncertainty in the geometric model; however, the quality of the grasp was not addressed, i.e., it was not a criterion for choosing the next evaluated point.

Gandler et al. addressed the problem so that probabilistic shape reconstruction from combined visual and tactile observations used sparse GPIS models [22]. They model the object surface as a sparse GP that is learned on the incomplete and noisy sensory data, and they explicitly represent uncertainty in different regions of the reconstruction. With the sparse GP, they reduced the computational cost, and in this way they enabled using larger visual and tactile datasets and more complex objects.

More recently, Chen et al. proposed a multi-finger sliding-touch strategy for exploring and reconstruction of unknown object shape from a partial depth observation and tactile measurements. BO is used to guide the exploration, while their single-leader strategy sets the fingers to get new information of the surface by sliding the contact [26]. The method focuses primarily on efficient geometric reconstruction rather than on selecting contacts according to grasp quality.

These approaches demonstrated how the GP or GPIS posterior uncertainty was used and was the basis for active surface reconstruction. However, their selection criteria, which

points should be evaluated next, are primarily designed to improve the geometric model. Some of the locations may be highly informative for shape reconstruction, without necessarily providing any information about the grasp. This distinction motivates approaches that consider surface information and grasp quality jointly.

3.2. Grasp Planning under Shape Uncertainty

A separate line of research considers the idea of how uncertainty in an object's shape can affect grasp planning. These methods used surface representations not only to describe the geometry of the object but also to evaluate whether a grasp is likely to remain stable across different possible realizations with the available sensory observations.

Mahler et al. proposed GP-GPIS-OPT, a grasp-planning method that uses a GPIS representation to account for shape uncertainty [11]. In their method, the Probability of Force Closure accounts for both shape uncertainty represented by the GPIS and uncertainty in physical grasp execution. Instead of evaluating force closure only on the posterior mean surface, the method considers a distribution over possible shapes that are consistent with the available observations. Although the Probability of Force Closure is formulated over possible object shapes and perturbed grasp configurations, direct evaluation would require computationally expensive integration over GPIS shape samples. For the tested objects, considering shape uncertainty yielded grasps with a higher probability of force closure than using a single deterministic surface estimate [11].

Unlike active tactile exploration, GP-GPIS-OPT plans parallel-jaw contact locations for a given GPIS model. It does not formulate the sequential acquisition of new geometric measurements or update the surface representation during grasp optimization. Furthermore, the presented formulation is limited to parallel-jaw grasps on two-dimensional object representations [11].

A more recent direction considers grasp transfer under partial geometric observation. De Farias et al. used GPIS-based reconstruction together with Functional Maps to transfer task-specific grasp functions from a reference object to a partially observed object that has the same category [27]. In contrast to active exploration approaches, the method relies on prior information from a related object and attempts to transfer a suitable grasp to it.

3.3. Bayesian Optimization for Robotic Grasping

Grasp evaluations can be expensive, noisy, and available only through a limited number of physical or simulated trials; therefore, Bayesian optimization is used to address this problem. A Gaussian process surrogate can model the relationship between grasp parameters and an observed grasp-quality measure, and an acquisition function is used to select the next grasp that should be evaluated [4, 23, 24, 25].

Nogueira et al. introduced Unscented Bayesian Optimization for safe robot grasping [23], where they addressed input uncertainty, such as errors in the relative pose between an object and the robot end effector. While the grasp can have a high objective value, it can remain unreliable if small execution errors produce a large reduction in grasp quality. To propagate input uncertainty through the Bayesian optimization decision process, they used unscented transformation. They favored grasp configurations that achieve high expected performance while remaining robust to local input perturbations. It was evaluated on synthetic optimization problems and realistic robot-grasp simulations.

Castanheira et al. extended Unscented Bayesian Optimization to three-dimensional haptic grasp exploration [24], where they combined the robustness properties of Unscented Bayesian Optimization with a collision-penalty heuristic. The method extends the search space from two to three translational hand-pose parameters, enabling the discovery of improved grasp configurations. When a collision occurs before finger closure, the collision-penalty heuristic assigns a penalized value to that configuration and guides subsequent evaluations away from the corresponding region. The approach was evaluated in grasping simulations [24].

More recently, Lechuz-Sierra et al. applied Bayesian optimization to robust robotic grasping using a sensorized compliant hand [25]. Their work used the tactile feedback to evaluate candidate grasps and investigated different grasp quality metrics within the BO. The method searches and optimizes the hand pose parameters to identify effective grasps for unknown objects in the presence of noise and uncertainty. It was evaluated both in simulation and on a physical robotic system equipped with a compliant, underactuated hand and tactile sensors [25].

Across these approaches, the Gaussian process surrogate models grasp quality as a function of the hand or grasp parameters rather than representing the object geometry as a probabilistic implicit surface. However, these approaches do not simultaneously update a GPIS-based surface reconstruction and optimize multiple individual three-dimensional

contact locations. Instead, they primarily optimize hand-pose or grasp parameters while treating the object geometry as fixed.

3.4. Joint Surface Exploration and Grasp Refinement

The work most closely related to the problem considered in this thesis is the method proposed by de Farias et al. [1]. Their approach simultaneously performs tactile exploration of an unknown object and refinement of a multi-finger grasp. An initial partial point cloud is acquired using an RGB-D camera and used to construct a GPIS model. During the iterations, the robot executes grasp attempts, obtains tactile observations, updates the GPIS, and evaluates grasp quality using Probability of Force Closure.

De Farias et al. used Bayesian optimization within a multi-dimensional hand configuration space. Their objective function integrates the Probability of Force Closure and incorporates a penalty term derived from the squared GPIS posterior mean evaluated at the resulting fingertip positions [1]:

$$y = \lambda P_{\mathrm{FC}} - \sum_{i=1}^{N_F} \overline{f}_{\mathrm{IS}} \left(\mathbf{x}_{F_i}\right)^2 , \tag{3.1}$$

where N_F denotes the number of fingers and $\overline{f}_{\mathrm{IS}}(\mathbf{x}_{F_i})$ is the GPIS predictive mean at the position of finger i [1]. Since the estimated surface is actually the zero level set of the GPIS mean, the second term penalizes fingertip locations for which the magnitude of the predictive mean is large. It therefore discourages configurations that are inconsistent with the current surface estimate. The scalar parameter λ, which is set empirically in the original work, scales the PFC contribution relative to the GPIS-based penalty [1, Sec. III-E].

This work and formulation are important because they connect three elements that had often been treated separately in earlier work: probabilistic surface reconstruction, tactile exploration, and probabilistic grasp evaluation. After each iteration, the GPIS is updated with contact data from grasp attempts, enabling the improved geometric model to guide future grasp searches. The reported experiments show that the approach can improve both the surface estimate and the number of stable grasps compared with a point-cloud-based sampling baseline [1]. In the formulation of de Farias et al., GPIS surface consistency penalty and PFC are combined into a single scalar target value.

Bayesian optimization updates a surrogate model of this combined target, and Expected Improvement is maximized to select the next query [1, Secs. III-B and III-E].

3.5. Research Gap

The reviewed literature establishes several important components of the problem reconstructing the object and finding a good grasp position. GP- and GPIS-based active exploration methods use geometric uncertainty to improve object reconstruction [10, 20, 21, 22]. Grasp planners such as GP-GPIS-OPT use probabilistic shape information to identify grasps with high Probability of Force Closure [11]. Bayesian optimization methods for robotic grasping improve sample efficiency and robustness under execution uncertainty [23, 24, 25]. De Farias et al. provide the closest joint framework by combining GPIS-based shape refinement with PFC-based grasp optimization [1].

However, none of the methods reviewed above uses a single acquisition function containing one component derived from the GPIS posterior and another component derived from the PFC surrogate to select joint contact configurations. Chen et al. [26] used BO to select tactile exploration actions for improving object-shape reconstruction. De Farias et al. [27] use GPIS-based reconstruction together with Functional Maps to transfer task-specific grasp functions from a reference object to a partially observed object of the same category. Among the reviewed methods, the earlier work and approach of de Farias et al. is the closest to the problem considered in this thesis. The method of de Farias et al. instead combines a GPIS-based surface-consistency penalty and PFC in a single scalar target, models this target with one Gaussian process surrogate, and applies Expected Improvement to select the next query [1, Secs. III-B and III-E].

To address this gap, this thesis proposes a geometry-aware composite acquisition function that combines information from the GPIS posterior and the PFC surrogate when selecting joint contact configurations. The GPIS-based component guides candidate contacts towards the estimated surface of the object while also favoring geometrically uncertain regions. The PFC-based Upper Confidence Bound favors joint configurations with either high predicted grasp quality or high uncertainty in that prediction. The combined acquisition function therefore guides the search towards geometrically plausible contact configurations while supporting both surface reconstruction and grasp-quality optimization to achieve higher evaluation scores for force closure.

4. Methodology

In this chapter, we want to present our approach, which incorporates together surface reconstruction of a partially observed object and the search for high-quality contacts grasp configurations in the acquisition function. We define the whole approach in three different sequential stages:

(i) a partial point cloud is obtained using a normal-based approximation of a one-sided scan with a predefined angular threshold;

(ii) a Gaussian Process Implicit Surface (GPIS) model is initially trained from the partial scan to encode a probabilistic representation of the object geometry; and

(iii) a Bayesian Optimization (BO) loop searches for configurations of N_c three-dimensional contact locations within the GPIS bounding box, guided jointly by the GPIS model and a Gaussian process surrogate of grasp quality.

4.1. Partial Point Cloud Acquisition

A partial point cloud is generated synthetically from a ground-truth triangular mesh of the object, where the points are obtained by simulating a one-sided scan with a predefined angle and using a normal-based approximation. Surface points are sampled uniformly from the mesh using area-weighted face sampling, and the corresponding outward face normal $\widehat{\mathbf{n}}_i$ is recorded for each sampled point $\mathbf{p}_i$.

To simulate sensor occlusion from a single scanner direction, only points whose outward normal is sufficiently aligned with the scanner direction are retained

$$\mathcal{P}_{\text{vis}} = \left\{ \mathbf{p}_i \;\middle|\; \widehat{\mathbf{n}}_i^{\top} \widehat{\mathbf{d}} > \cos\theta_{\max} \right\}, \tag{4.1}$$

where $\mathcal{P}_{\text{vis}}$ denotes the set of retained surface points, $\mathbf{p}_i$ is the i-th sampled point, and $\widehat{\mathbf{n}}_i$ is its unit outward normal. The unit vector $\widehat{\mathbf{d}}$ defines the scanner direction, while $\theta_{\max}$ is the maximum accepted angle between the surface normal and the scanner direction. Since both vectors are normalized, their dot product satisfies $\widehat{\mathbf{n}}_i^{\top}\widehat{\mathbf{d}} = \cos\theta_i$, where θ_i is the angle between them. Therefore, the condition in Equation 4.1 retains a point only when $\theta_i < \theta_{\max}$. A smaller value of $\theta_{\max}$ produces a narrower partial scan, and a larger value retains a greater portion of the object surface.

We simulate the sensor noise by adding independent zero-mean Gaussian perturbations to each saved point $\widetilde{\mathbf{p}}_i = \mathbf{p}_i + \varepsilon_i$ where $\varepsilon_i \sim \mathcal{N}\left(\mathbf{0}, \sigma_s^2 I_3\right)$. The resulting point cloud $\widetilde{\mathcal{P}}_{\text{vis}}$, together with the associated surface normals, is stored and used as the input to the next phase, GPIS initial training stage.

The transition from the full surface point cloud to the partially observed point cloud is illustrated in the left part of Figure 4.1.

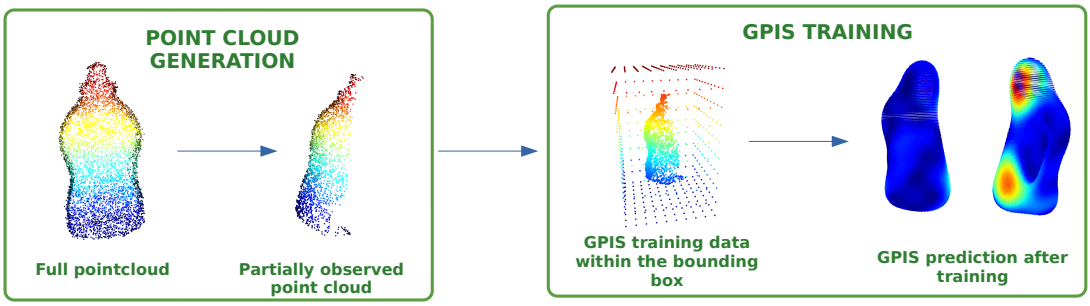


Figure 4.1.: Overview of partial point cloud generation and initial GPIS training. On the left, a complete surface point cloud sampled from the input mesh is shown for reference, while the partially observed point cloud is obtained by filtering the sampled points using the predefined angular threshold. On the right, the resulting partial point cloud is placed inside an expanded bounding box, whose sampled surface points provide outside observations for GPIS training. The two views on the far right show the posterior prediction of the trained GPIS model from opposite viewpoints.

4.2. Gaussian Process Implicit Surface

The partial point cloud described in Section 4.1 is used to initialize a GPIS model of the object geometry. The model provides both an estimate of the implicit surface and uncertainty information, which are used during the subsequent stage of reconstructing the object and searching for a good grasp position. As illustrated in the right part of Figure 4.1, the partial point cloud is enclosed by an expanded bounding box and used to construct the initial GPIS training set. Points generated on the bounding box, which provide initial outside constraints, are labeled as outside observations

4.2.1. Implicit Surface Representation

A Gaussian Process Implicit Surface (GPIS) uses a Gaussian process to represent the object surface through an implicit function $\phi : \mathbb{R}^3 \to \mathbb{R}$, where surface points correspond to $\phi(\mathbf{x}) = 0$. The sign convention adopted here is

$$\phi(\mathbf{x}) = \begin{cases} 0 & \text{on the object surface,} \\ +1 & \text{outside the object,} \\ -1 & \text{inside the object.} \end{cases} \tag{4.2}$$

The inward surface normal at any surface point is estimated from the negative gradient of the GPIS posterior mean and normalized to unit length. Therefore, we label two categories of points that are assembled from the partial point cloud:

1. **Surface points** ($\phi = 0$): the scanned surface points $\widetilde{\mathbf{p}}_i \in \widetilde{\mathcal{P}}_{\text{vis}}$.
2. **Outside points** ($\phi = +1$): points sampled on the surface of an expanded axis-aligned bounding box constructed from the ground-truth object-mesh bounds.

The initial GPIS training set is

$$\mathcal{D}_\phi = \{(\mathbf{x}_j, y_j)\}_{j=1}^{N_\phi}, \qquad \mathbf{x}_j \in \mathbb{R}^3, \quad y_j \in \{0, +1\}.$$

4.2.2. Sparse GPIS

The inference process for full Gaussian Processes (GP) or Gaussian Process Implicit Surfaces (GPIS) becomes significantly slower as the number of training points increases, resulting in high computational costs for large point clouds. Therefore, the GPIS is implemented using a sparse Gaussian process approximation with M inducing points [28]. The inducing locations are initialized by randomly selecting M points from the GPIS training set.

The prior over the latent function is defined as

$$\phi \sim \mathcal{GP}\big(0,\ k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta})\big)\,, \tag{4.3}$$

where k is a Matern kernel [2] with $\nu = 5/2$ and learnable hyperparameters $\boldsymbol{\theta}$.

Observations are modeled with additive Gaussian noise

$$y_j = \phi(\mathbf{x}_j) + \varepsilon_j^n, \qquad \varepsilon_j^n \sim \mathcal{N}(0, \sigma_n^2). \tag{4.4}$$

The parameters of the sparse GP model are optimized using the Adam optimizer by minimizing a loss based on the evidence lower bound (ELBO) [28]. Early stopping is applied when the training loss no longer improves.

Given a query point $\mathbf{x}^* \in \mathbb{R}^3$, the GPIS posterior predictive distribution is Gaussian with mean $\mu_\phi(\mathbf{x}^*)$ and standard deviation $\sigma_\phi(\mathbf{x}^*)$. The probability that the implicit function value lies within a tolerance interval around the surface is

$$p_{\text{surf}}(\mathbf{x}^*) = \Phi\left(\frac{\varepsilon - \mu_\phi(\mathbf{x}^*)}{\sigma_\phi(\mathbf{x}^*)}\right) - \Phi\left(\frac{-\varepsilon - \mu_\phi(\mathbf{x}^*)}{\sigma_\phi(\mathbf{x}^*)}\right), \tag{4.5}$$

where Φ denotes the standard normal cumulative distribution function (CDF) and $\varepsilon > 0$ is the surface-proximity tolerance. In other words, it is equivalent to the posterior probability that $|\phi(\mathbf{x}^*)| \leq \varepsilon$, i.e. point on the surface is.

4.3. Geometry-Aware Bayesian Optimization for Multi-Contact Grasp Search

After the initial GPIS model is constructed, Bayesian optimization is used to search for contact configurations such that are both geometrically plausible and likely to achieve a high probability in the force closure metric. The search is guided by the initial GPIS estimate of the object surface and by a separate Gaussian process surrogate trained on evaluated probability-of-force-closure values. These two sources of information - their outputs are combined within the acquisition function described in the following subsections. Both models are updated after each iteration.

4.3.1. Problem Formulation

The objective is to find the configuration of N_c contact points on the object surface that provides a high probability of force closure. During the search, we also explore and reconstruct the object. As previously unobserved surface regions are incorporated into the GPIS model, the surface of the reconstructed object is expanded, thereby increasing the number of contact configurations that may provide a high probability of force closure.

Each contact position is bounded by the bounding box of the GPIS model. Therefore, the joint search domain $\mathcal{X} \subset \mathbb{R}^{3N_c}$ is obtained by applying these three-dimensional bounds independently to all N_c contacts.

Evaluating a candidate configuration $\mathbf{x}$ requires querying the GPIS and running a stochastic force-closure evaluation using Monte Carlo sampling. Bayesian optimization is used to find a high-quality grasp position within a limited evaluation budget.

4.3.2. GP Surrogate for Probability of Force Closure

A Gaussian process surrogate is used to predict the probability of force closure for candidate contact configurations. At iteration t, the surrogate is trained on the previously evaluated configurations:

$$\mathcal{D}_{\mathrm{PFC},t} = \left\{ \left(\mathbf{x}_j, \widehat{P}_{\mathrm{FC},j} \right) \right\}_{j=1}^{t}, \tag{4.6}$$

where $\mathbf{x}_j$ is an evaluated contact configuration and $\widehat{P}_{\mathrm{FC},j} \in [0,1]$ is its Monte Carlo estimate of the probability of force closure. The surrogate uses a Matérn kernel with $\nu = 5/2$. Its posterior mean $\mu_{\mathrm{PFC},t}(\mathbf{x})$ and standard deviation $\sigma_{\mathrm{PFC},t}(\mathbf{x})$ are updated after each evaluation.

4.3.3. Warm-up Phase

At the beginning of the BO, the PFC surrogate initially contains no evaluated contact configurations. Therefore, optimization starts with a warm-up phase. During this phase, different contact configurations are sampled uniformly within the bounding box. Each of the sampled configurations is evaluated through the initial GPIS model and Monte Carlo force-closure procedure. After each iteration, geometric observations, i.e., each point is added to the GPIS model, while the resulting PFC estimate is added to the PFC surrogate dataset. After the warm-up phase is done, the PFC surrogate model contains an initial set of observations. After the warm-up phase, candidate configurations are selected using the acquisition function. However, the exploration parameter remains fixed at $\kappa_{\max}$ until the initial exploration period is completed.

4.3.4. GPIS-Guided UCB Acquisition Function

The acquisition function combines information from the GPIS model and the PFC surrogate. GPIS part goes for contact configurations that are close to the estimated object surface or lie in geometrically uncertain regions. The PFC component prefers configurations with a high predicted probability of force closure or high surrogate uncertainty. A diversity bonus is additionally included to discourage contacts from clustering at the same location. The construction of the composite acquisition function is illustrated in Figure 4.2.

For a candidate joint contact configuration $\mathbf{x} = [\mathbf{c}_1^\top, \ldots, \mathbf{c}_{N_c}^\top]^\top$, where N_c denotes the number of contacts and $\mathbf{c}_i \in \mathbb{R}^3$ is the position of the i-th contact, the GPIS-derived acquisition component is defined as

$$\alpha_{\mathrm{GPIS},t}(\mathbf{x}) = \frac{1}{s_{\mathrm{GPIS}}} \left[-\frac{1}{N_c} \sum_{i=1}^{N_c} |\mu_{\phi,t}(\mathbf{c}_i)| + \kappa_t \frac{1}{N_c} \sum_{i=1}^{N_c} \sigma_{\phi,t}(\mathbf{c}_i) \right], \tag{4.7}$$

where $\mu_{\phi,t}(\mathbf{c}_i)$ is the GPIS posterior mean and $\sigma_{\phi,t}(\mathbf{c}_i)$ is the corresponding posterior standard deviation at contact location $\mathbf{c}_i$ and iteration t. The scale factor $s_{\mathrm{GPIS}} > 0$ is

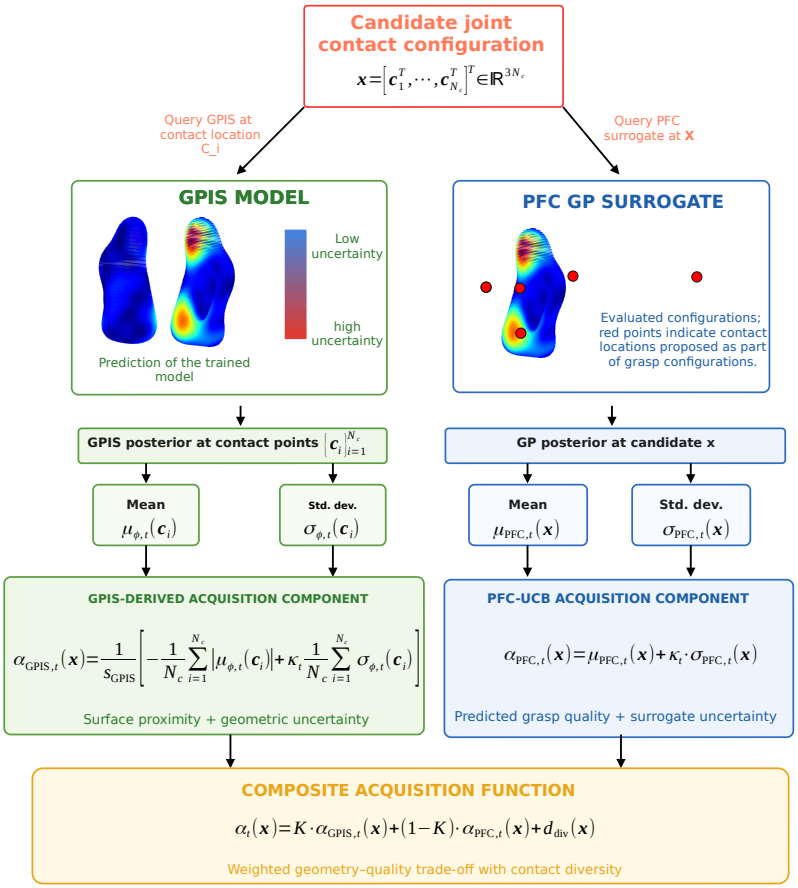


Figure 4.2.: Construction of the geometry-aware composite acquisition function. A candidate joint contact configuration $\mathbf{x}$ is decomposed into individual contact locations, which are queried using the GPIS model to obtain the posterior mean and standard deviation. The PFC GP surrogate is queried directly at the full configuration $\mathbf{x}$. The resulting GPIS-derived and PFC-UCB components are combined using the weighting factor K, together with the contact diversity bonus, to obtain the composite acquisition value. Here, s_{GPIS} is a scale factor used to normalize the GPIS-derived acquisition component, while κ_t is the iteration-dependent exploration parameter that controls the contribution of posterior uncertainty in both acquisition components.

computed once at initialization as the standard deviation of the average GPIS proximity values obtained from a set of randomly sampled contact configurations. The factor $1/s_{\mathrm{GPIS}}$ normalizes the GPIS-derived acquisition component before it is combined with the PFC-based component. The first term favors contact locations whose predicted implicit-function values are close to zero and are therefore close to the estimated object surface. The second term favors contact locations with high GPIS posterior uncertainty and therefore promotes geometric exploration. The parameter $\kappa_t \geq 0$ controls the contribution of this exploration term. Both terms are averaged over the N_c contacts so that the magnitude of the acquisition component does not increase simply because more contacts are used.

In parallel, the PFC component is obtained from the posterior of the separate Gaussian process surrogate:

$$\alpha_{\mathrm{PFC},t}(\mathbf{x}) = \mu_{\mathrm{PFC},t}(\mathbf{x}) + \kappa_t \sigma_{\mathrm{PFC},t}(\mathbf{x}), \tag{4.8}$$

where $\mu_{\mathrm{PFC},t}(\mathbf{x})$ is the predicted probability of force closure and $\sigma_{\mathrm{PFC},t}(\mathbf{x})$ represents the surrogate uncertainty. The PFC acquisition value is high when the surrogate predicts a high probability of force closure, has high uncertainty, or both.

To discourage the contact points from clustering, a diversity bonus is defined using the minimum pairwise distance between contacts:

$$d_{\mathrm{div}}(\mathbf{x}) = \lambda_{\mathrm{div}} \min_{i \neq j} \|\mathbf{c}_i - \mathbf{c}_j\|_2 , \tag{4.9}$$

where $\lambda_{\mathrm{div}} \geq 0$ determines the contribution of the diversity term. Maximizing this term encourages the closest pair of contacts to remain separated.

The GPIS-derived component, the PFC-UCB component, and the diversity bonus are combined into the final acquisition function as

$$\alpha_t(\mathbf{x}) = K\, \alpha_{\mathrm{GPIS},t}(\mathbf{x}) + (1 - K)\, \alpha_{\mathrm{PFC},t}(\mathbf{x}) + d_{\mathrm{div}}(\mathbf{x}), \tag{4.10}$$

where $K \in [0, 1]$ controls the relative contribution of geometry-based guidance and predicted grasp quality.

4.3.5. Acquisition Optimization

At each post-warm-up iteration, a set of random candidate configurations is sampled uniformly and evaluated using Equation (4.10). A certain number of candidates with the highest scores are used as starting points for independent L-BFGS optimization runs [29].

Depending on the selected restart strategy, a recently evaluated configuration or the best previously evaluated configuration may also be included as an additional starting point. Each run minimizes $-\alpha_t(\mathbf{x})$ using L-BFGS with a strong-Wolfe line search, which selects the step length such that sufficient-decrease and curvature conditions are satisfied [29]. During optimization, the contact coordinates are transformed so that they always remain within the bounding box of the GPIS model. From the group of refined candidates, we select the one with the highest acquisition value for evaluation.

4.3.6. Candidate Evaluation

After optimizing the acquisition function, a candidate configuration $\mathbf{x}_t$ is selected and its contacts are first evaluated using the GPIS model. The surface cost is defined as

$$s_{\mathrm{surf},t}(\mathbf{x}_t) = 1 - \frac{1}{N_c}\sum_{i=1}^{N_c} p_{\mathrm{surf},t}(\mathbf{c}_{i,t}), \tag{4.11}$$

where $p_{\mathrm{surf},t}(\mathbf{c}_{i,t})$, defined in Equation (4.5), is the posterior probability that the i-th contact lies on the surface or within the selected tolerance around the estimated surface. A lower surface cost indicates that the contact configuration is more consistent with the current GPIS surface estimate.

The stochastic force-closure evaluation is performed only for candidate configurations whose surface cost is below a predefined threshold τ_{gate}

$$s_{\mathrm{surf},t}(\mathbf{x}_t) \leq \tau_{\mathrm{gate}}. \tag{4.12}$$

Here, τ_{gate} is the maximum surface cost allowed before the Monte Carlo force-closure evaluation is performed. If the surface cost is higher than this threshold, we discard the candidate and classify it as a bad choice w.r.t the current GPIS surface estimate, and set PFC to zero for this configuration of contacts.

For each perturbed grasp, friction cones are represented by a finite set of primitive contact forces, which are then used to construct the corresponding contact wrenches. The Ferrari–Canny ε-metric is computed from the convex hull of these wrenches [19]. If $\varepsilon^{(k)}(\mathbf{x}_t)$ is positive, it means the perturbed grasp achieves force closure. For Monte Carlo trial k, the binary outcome is defined as

$$Y_k = \begin{cases} 1, & \text{if } \varepsilon^{(k)}(\mathbf{x}_t) > 0, \\ 0, & \text{otherwise.} \end{cases} \tag{4.13}$$

The probability of force closure is estimated as

$$\widehat{P}_{\mathrm{FC}}(\mathbf{x}_t) = \frac{1}{N_{\mathrm{MC}}} \sum_{k=1}^{N_{\mathrm{MC}}} Y_k, \tag{4.14}$$

where N_{MC} is the number of Monte Carlo trials, and $\varepsilon^{(k)}(\mathbf{x}_t)$ is the Ferrari–Canny metric obtained in trial k. Therefore, we define $\widehat{P}_{\mathrm{FC}}(\mathbf{x}_t)$ as the fraction of Monte Carlo trials in which the perturbed grasp achieves force closure.

A combined score is used to monitor the search and identify the best evaluated contact configuration, denoted by $\mathbf{q}_{\mathrm{best}}$,

$$q_{\mathrm{comb},t}(\mathbf{x}_t) = K \left(-\frac{s_{\mathrm{surf},t}(\mathbf{x}_t)}{s_{\mathrm{GPIS}}} \right) + (1 - K)\widehat{P}_{\mathrm{FC}}(\mathbf{x}_t), \tag{4.15}$$

where $q_{\mathrm{comb},t}(\mathbf{x}_t)$ is the combined score of the evaluated configuration $\mathbf{x}_t$ at iteration t. The parameter $K \in [0, 1]$ controls the relative contribution of the surface cost and the probability of force closure. A larger value of K gives more importance to how close the contacts are to the estimated object surface, while a smaller value gives more importance to the evaluated PFC value. The term $s_{\mathrm{surf},t}(\mathbf{x}_t)$ is the surface cost defined in Equation (4.11). Since a lower surface cost is better, it is multiplied by -1, so that configurations with lower surface cost obtain a larger contribution to the combined score. The scale factor s_{GPIS} is the same factor used in the GPIS acquisition component. It places the surface-related term and the PFC term on more comparable scales before they are combined. The value $\widehat{P}_{\mathrm{FC}}(\mathbf{x}_t)$ is the Monte Carlo estimate of the probability of force closure defined in Equation (4.14). A higher PFC value increases the combined score.

This score is used only for monitoring the search and selecting $\mathbf{q}_{\mathrm{best}}$, the evaluated configuration with the highest combined score. The PFC surrogate is trained directly on the evaluated PFC values and not on the combined score.

4.3.7. Adaptive Exploration Parameter

The exploration parameter is initially fixed at its maximum value $\kappa_{\max}$. During the warm-up phase, as already mentioned in Section 4.3.3, configurations are sampled uniformly at random. After the warm-up phase, the acquisition function is used to select candidate configurations, while κ_t remains fixed at $\kappa_{\max}$ until a certain number of evaluations have

been completed. After this initial exploration period, κ_t is adapted using the best surface cost and the highest PFC estimate observed since the adaptive phase began.

The surface-related progress signal is defined as

$$r_{s,t} = 1 - \text{clip}\left(\frac{s^{\min}_{\text{surf},t}}{\tau_s}, 0, 1\right), \tag{4.16}$$

where $s^{\min}_{\text{surf},t}$ is the lowest surface cost observed since the adaptive phase began, defined in Equation 4.11, and τ_s is an empirically selected threshold used to measure geometric progress. The signal $r_{s,t}$ remains equal to zero while the lowest observed surface cost is greater than or equal to τ_s. In the moment when the surface cost falls below this threshold, a further decrease in the surface cost results in a larger value of $r_{s,t}$.

The PFC-related progress signal is defined as

$$r_{p,t} = \text{clip}\left(\frac{\widehat{P}^{\max}_{\text{FC},t}}{\tau_p}, 0, 1\right), \tag{4.17}$$

where $\widehat{P}^{\max}_{\text{FC},t}$ is the highest PFC estimate observed since the adaptive phase began, and τ_p is the target PFC value empirically selected. The signal $r_{p,t}$ increases as higher PFC values are found and reaches one when the highest observed PFC estimate is equal to or greater than τ_p.

The clipping operation limits both progress signals to the interval $[0, 1]$. The two progress signals are combined as

$$\text{combined}_t = (r_{s,t} \cdot r_{p,t})^4 . \tag{4.18}$$

The two signals are multiplied so that the combined progress is high only when both the surface cost and the PFC estimate show good results. Using the fourth power, i.e., raising their product to the fourth power, keeps the combined progress low until both signals are close to one. This delays the correction of κ_t and keeps the search exploratory for a longer period.

The exploration parameter is updated according to

$$\kappa_t = \kappa_{\min} + (\kappa_{\max} - \kappa_{\min})\,(1 - \text{combined}_t) . \tag{4.19}$$

In the case when one or both progress signals stay low, κ_t remains close to $\kappa_{\max}$, and gives greater importance to exploration. As both progress signals approach their target values, κ_t decreases towards $\kappa_{\min}$.

4.3.8. Overall Optimization Procedure

Algorithm 1 summarizes the complete optimization procedure. The complete optimization procedure begins with the warm-up phase, during which contact configurations are sampled uniformly at random. For each configuration, we evaluate using the current GPIS model and the Monte Carlo force-closure procedure. The resulting PFC is added to the PFC surrogate dataset, while the geometric observations obtained at the contact locations are added to the GPIS training data.

After the warm-up phase, candidate configurations are selected by optimizing the acquisition function defined in Equation (4.10). After one configuration is selected, first we check it using the surface-cost criterion. Only if the cost is below of the predefined surface threshold for PFC, the probability of force closure will be estimated using Monte Carlo sampling. Otherwise, its PFC value is set to zero.

Following each evaluation, the observation $(\mathbf{x}_t, \widehat{P}_{\mathrm{FC}}(\mathbf{x}_t))$ is added to the PFC surrogate dataset, and the geometric information obtained from the evaluated contacts is added to the GPIS training data. The posterior predictions of both models are then updated. Once the initial exploration is complete, the exploration parameter κ_t is updated using the best surface cost and the highest PFC estimate observed during the adaptive phase, as explained in Section 4.3.7. The acquisition function is then optimized again using the updated models and exploration parameter.

This sequence of candidate selection, evaluation, model update, and adaptation is repeated until the maximum number of evaluations is reached. The evaluated configuration with the highest combined score $\mathbf{q}_{\mathrm{best}}$, is returned as the final grasp configuration.

Algorithm 1: Geometry-Aware Bayesian Optimization for Multi-Contact Grasp Search

Input: Initial GPIS model, Gaussian-process surrogate for PFC, search domain $\mathcal{X}$, maximum number of evaluations T, number of warm-up evaluations N_{warm}, start of adaptive exploration N_{adapt}, exploration bounds $\kappa_{\min}$ and $\kappa_{\max}$, surface-progress threshold τ_s, target PFC value τ_p, and surface-cost gating threshold τ_{gate}

Output: Contact configuration with the highest combined score $\mathbf{q}_{\text{best}}$

$\mathcal{D}_{\text{PFC}} \leftarrow \emptyset$;
$\kappa_1 \leftarrow \kappa_{\max}$;
$q_{\text{best}}^{\text{score}} \leftarrow -\infty$;
$\mathbf{q}_{\text{best}} \leftarrow \varnothing$;
for $t \leftarrow 1$ **to** T **do**
  if $t \leq N_{\text{warm}}$ **then**
    Sample $\mathbf{x}_t$ uniformly from $\mathcal{X}$;
  else
    Select $\mathbf{x}_t$ by maximizing $\alpha_t(\mathbf{x})$ from Equation (4.10);
  end
  Compute $s_{\text{surf},t}(\mathbf{x}_t)$ using Equation (4.11);
  if $s_{\text{surf},t}(\mathbf{x}_t) \leq \tau_{\text{gate}}$ **then**
    Estimate $\widehat{P}_{\text{FC}}(\mathbf{x}_t)$ using Monte Carlo sampling;
  else
    $\widehat{P}_{\text{FC}}(\mathbf{x}_t) \leftarrow 0$;
  end
  Obtain geometric observations at the evaluated contacts;
  Add the geometric observations to the GPIS training data;
  Update the GPIS posterior;
  Add $(\mathbf{x}_t, \widehat{P}_{\text{FC}}(\mathbf{x}_t))$ to $\mathcal{D}_{\text{PFC}}$;
  Update the PFC surrogate posterior;
  Compute $q_{\text{comb},t}(\mathbf{x}_t)$ using Equation (4.15);
  if $q_{\text{comb},t}(\mathbf{x}_t) > q_{\text{best}}^{\text{score}}$ **then**
    $q_{\text{best}}^{\text{score}} \leftarrow q_{\text{comb},t}(\mathbf{x}_t)$;
    $\mathbf{q}_{\text{best}} \leftarrow \mathbf{x}_t$;
  end
  if $t \geq N_{\text{adapt}}$ **then**
    Compute $r_{s,t}$ using Equation (4.16);
    Compute $r_{p,t}$ using Equation (4.17);
    Update κ_{t+1} using Equation (4.19);
  else
    $\kappa_{t+1} \leftarrow \kappa_{\max}$;
  end
end
return $\mathbf{q}_{\text{best}}$;

5. Experiments

This chapter presents an evaluation of the proposed approach, referred to as the *Geometry-Aware Acquisition Function*, within a Bayesian optimization framework for object reconstruction and the search for qualitative grasp configurations. The contact-search problem considered in this thesis is especially challenging due to the independent optimization of four three-dimensional contact positions, which results in a 12-dimensional Cartesian search space. In the experiments, the main focus is on the composite acquisition function, while the additional acquisition strategies are used as baselines. The experiments analyze the effects of the exploration parameter κ, the L-BFGS restart strategy, and object geometry. The geometry-aware acquisition function is subsequently compared with other acquisition functions, such as vanilla UCB, Expected Improvement, and a method based on the combined GPIS–PFC formulation of de Farias et al. [1]. The selected contact configurations are evaluated quantitatively using the GPIS surface cost and the estimated Probability of Force Closure (PFC). Interpretation of the observed behavior and the limitations of the evaluation are discussed in Chapter 6.

5.1. Experimental Setup

This section describes the common experimental setup used throughout the evaluation. It introduces the test objects, the main implementation details, and the parameters used throughout the experiments. The experimental framework was implemented in Python, with Pyro used for the Gaussian-process models and Bayesian optimization [30].

5.1.1. Test Objects

Experiments are conducted on five object models from the YCB Object and Model Set [31]: the mustard bottle, tuna fish can, cracker box, mini soccer ball, and Rubik's cube. In the

YCB set, these are `006_mustard_bottle`, `007_tuna_fish_can`, `003_cracker_box`, `053_mini_soccer_ball`, and `077_rubiks_cube`. The objects differ in shape and surface structure, from simple symmetric forms to more irregular and elongated ones. This selection enables evaluation of the contact-search method across diverse object geometries.

5.1.2. Initial Partial Observations

For each object, a partial point cloud is generated from the corresponding ground-truth mesh using the procedure described in Section 4.1. A total of $5\,000$ surface points are sampled from the mesh. The scanner half-angle is set to $\theta_{\max} = 80^\circ$, and a fixed scanner direction is selected for each object such that a substantial part of the surface is observed while some of the geometry remains unobserved. The scanner directions used for the individual objects are listed in Table 5.1. The results of the partial-observation procedure are illustrated in Figure 5.1, where the complete surface point cloud of each object is shown together with the corresponding initial partial observation.

Table 5.1.: Scanner directions used to generate the initial partial observation for each object.

Object	**Scanner direction**
Mustard bottle	$(1,\ 0,\ 0)$
Tuna fish can	$(0.5,\ 0,\ 1)$
Cracker box	$(1,\ 1,\ 0)$
Mini soccer ball	$(1,\ 0,\ 0)$
Rubik's cube	$(1,\ 0,\ 0)$

Gaussian sensor noise with a standard deviation of $\sigma_s = 1\,\mathrm{mm}$ is added to the sampled points. The resulting partial point cloud and corresponding surface normals define the initial geometric observation used to construct the GPIS model. Partial observations are used because the proposed method is intended to find high-quality grasp configurations while reconstructing an initially unknown object. GPIS models provide a probabilistic representation of incomplete object geometry and retain uncertainty in regions that have not been sufficiently observed [10, 22].

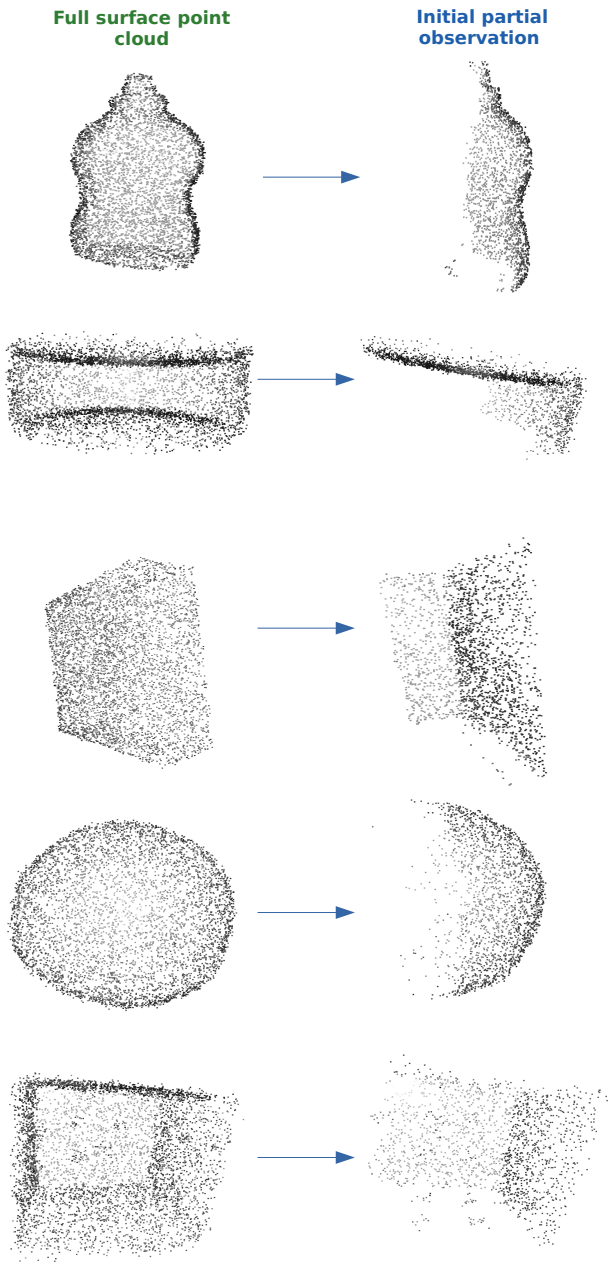


Figure 5.1.: Complete surface point clouds and the corresponding initial partial observations for the five test objects. From top to bottom: mustard bottle, tuna fish can, cracker box, mini soccer ball, and Rubik's cube.

5.1.3. Initial GPIS Models

The initial GPIS model for each object is trained from the corresponding partial point cloud using the procedure described in Section 4.2. The observed surface points are assigned an implicit value of zero. The bounding box is constructed from the axis-aligned bounds of the ground-truth object mesh and expanded by $0.02\,\mathrm{m}$ in each coordinate direction. Outside observations with implicit value $+1$ are generated on the six faces of this bounding box using a 10×10 grid on each face. A sparse Gaussian process with $M = 150$ inducing points and a Matérn-$5/2$ covariance function is used. The Matérn family is a standard Gaussian-process covariance family and is also included in the GP-UCB analysis of Srinivas et al. [2, 7]. The observation-noise parameter is initialized to $\sigma_n = 0.003$. Initial GPIS training is performed for at most $2\,000$ Adam steps at a learning rate of 10^{-2}. Early stopping is triggered if the loss does not improve by more than 10^{-4} for 70 consecutive optimization steps. The resulting GPIS model provides the initial posterior mean and uncertainty before Bayesian optimization begins.

During the subsequent contact-search loop, the initial GPIS is not treated as a complete object model. Each evaluated joint contact configuration provides new simulated geometric observations, which are appended to the GPIS data. The GPIS posterior used in later acquisition evaluations therefore reflects the additional information obtained during the search.

5.1.4. Bayesian Optimization Protocol

Each candidate grasp is represented by four three-dimensional contact locations, resulting in a 12-dimensional Cartesian search space,

$$\mathbf{x} = \left[\mathbf{c}_1^T, \mathbf{c}_2^T, \mathbf{c}_3^T, \mathbf{c}_4^T\right]^T \in \mathbb{R}^{12}. \tag{5.1}$$

The search domain is the expanded axis-aligned bounding box used by the GPIS model, with the same three-dimensional bounds applied independently to each of the four contact locations. The same contact-space parameterization and search domain are used for the geometry-aware acquisition function and all baseline strategies.

Each standard run consists of 350 evaluations. The first 50 evaluations serve as a random warm-up phase, during which joint contact configurations are uniformly sampled from the search domain. Each warm-up evaluation is used to update both the GPIS model and the initially empty PFC surrogate. Acquisition-guided selection begins after the warm-up

phase. For configurations using adaptive κ, the value is kept at $\kappa_{\max}$ until evaluation 100, after which the adaptive update is enabled.

At each acquisition-guided iteration, $2\,000$ joint configurations are sampled uniformly from the search domain and evaluated by the corresponding acquisition function. The three highest-scoring candidates are used as L-BFGS starting points. A fourth starting point is supplied according to the restart strategy described in Section 4.3.5. L-BFGS uses a strong-Wolfe line search and a maximum of 200 inner iterations.

The PFC Gaussian-process surrogate is initialized without prior PFC training data. A Matérn-$5/2$ kernel is used, and every evaluated pair $(\mathbf{x}_t, \widehat{P}_{\mathrm{FC}}(\mathbf{x}_t))$ is added to the surrogate dataset.

After a joint contact configuration has been selected, the ground-truth mesh is used to simulate the corresponding geometric observation. A selected point that lies inside the object or within $0.01\,\mathrm{m}$ of the mesh is projected onto the closest surface point and added to the GPIS data with implicit value 0. Points outside this region are added with implicit value $+1$. The GPIS posterior is subsequently updated and used during later acquisition evaluations.

Grasp quality is evaluated using the Probability of Force Closure described in Section 2.4.2. In the experiments, Monte Carlo PFC evaluation is performed only when the GPIS surface cost satisfies

$$s_{\mathrm{surf}} \leq \tau_{\mathrm{gate}}, \qquad \tau_{\mathrm{gate}} = 0.1. \tag{5.2}$$

Otherwise, the candidate is assigned a PFC value of zero. When the gate is satisfied, $K_{\mathrm{mc}} = 10$ Monte Carlo trials are used. Contact positions are perturbed with a standard deviation of $10\,\mathrm{mm}$, surface normals with a standard deviation of 0.01, and the center of mass with a standard deviation of $5\,\mathrm{mm}$. The friction coefficient is sampled from $\mathcal{N}(0.2, 0.05^2)$ and clipped to $[0.2, 0.5]$.

All standard configurations are evaluated on all five objects using five random seeds, $\{42, 2, 5, 7, 17\}$, resulting in 25 runs for each acquisition and hyperparameter configuration.

Table 5.2 summarizes the fixed experimental parameters. Acquisition-specific settings are reported separately in Table 5.3.

Table 5.2.: Fixed parameters used in the standard experiments.

Component	Parameter	Value
Point cloud	sampled surface points	$5\,000$
	scanner half-angle	80°
	sensor-noise std.	$1\,\mathrm{mm}$
GPIS	kernel	Matérn-$5/2$
	inducing points	150
	bounding-box margin	$0.02\,\mathrm{m}$
	observation noise	0.003
	surface tolerance ε	0.10
	maximum training steps	$2\,000$
	Adam learning rate	10^{-2}
Bayesian optimization	number of contacts	4
	search dimension	12
	total evaluations	350
	standard warm-up	50
	random candidates per iteration	$2\,000$
	L-BFGS random starts	3
	additional L-BFGS start	1
	L-BFGS maximum inner iterations	200
	adaptive-κ start	100
	objects	5
	random seeds	$42, 2, 5, 7, 17$
PFC evaluation	surface gate τ_{gate}	0.1
	Monte Carlo trials	10
	contact-position std.	$10\,\mathrm{mm}$
	normal perturbation std.	0.01
	center-of-mass std.	$5\,\mathrm{mm}$
	friction coefficient	$\mathcal{N}(0.2, 0.05^2)$, clipped to $[0.2, 0.5]$

5.2. Geometry-Aware Acquisition Function and Baselines

The Geometry-Aware Acquisition Function introduced in Chapter 4 is the primary method evaluated in this chapter. Three additional acquisition strategies are used as baselines to show the contribution of the geometry-aware acquisition function.

Our method combines the GPIS-derived and PFC-based acquisition components. Its GPIS component prefers contacts close to the estimated zero-level surface while rewarding geometric uncertainty, whereas the PFC component uses

$$\mu_{\mathrm{PFC}} + \kappa\sigma_{\mathrm{PFC}}.$$

The two components are combined with $K = 0.5$, together with a diversity bonus of weight 0.2. Three exploration settings are evaluated for the Geometry-Aware Acquisition Function: adaptive $\kappa : 4.0 \to 0.8$, fixed $\kappa = 2$, and fixed $\kappa = 3$. For the adaptive setting, $\kappa = 4.0$ is retained through the first 100 evaluations and is subsequently adapted according to the rule introduced in Chapter 4, with a lower bound of $\kappa_{\min} = 0.8$. The fixed settings use $\kappa = 2$ or $\kappa = 3$ throughout the acquisition-guided phase. We also evaluated our approach using the two L-BFGS restart strategies introduced in Chapter 4: `last` and `best`.

The first baseline, PFC-only UCB, removes the GPIS-derived acquisition component and selects configurations using only UCB on the PFC surrogate. Fixed values $\kappa \in \{0, 2, 3\}$ are evaluated, with all runs using the `last` restart mode. The second baseline uses Expected Improvement to the PFC surrogate. The values $\xi \in \{0.0, 0.1, 0.2\}$ are evaluated, and all runs use the `last` restart mode. The third baseline follows the combined-target structure of de Farias et al., where they combined the GPIS-derived geometric signal and the PFC into a single scalar target and modeled it using a single Gaussian-process surrogate. The method is evaluated with adaptive $\kappa : 4.0 \to 0.8$, and fixed $\kappa = 2$, using the `last` restart mode. Table 5.3 summarizes the tested acquisition configurations.

5.2.1. Hand-Inspired Contact Configuration Experiment

In addition to the standard experiments, we tested the geometry-aware acquisition function method using a simplified, hand-inspired contact setup. This experiment uses the same four-contact, 12-dimensional search space as the standard experiments described in Section 5.1.4, and each contact is represented by a three-dimensional position $\mathbf{c}_i \in \mathbb{R}^3$.

Table 5.3.: Acquisition strategies and parameter settings used in the standard experiments.

Method	Parameter setting	Restart mode
Geometry-Aware Acquisition Function	$4.0 \rightarrow 0.8,\ 2,\ 3$	`last`, `best`
PFC-only UCB	$\kappa = 0,\ 2,\ 3$	`last`
Expected Improvement	$\xi = 0.0,\ 0.1,\ 0.2$	`last`
Combined-target baseline based on de Farias et al.	$4.0 \rightarrow 0.8,\ 2$	`last`

Our intuition is to see whether the geometry-aware acquisition still works well when the four contact points are encouraged to form a spatially constrained, hand-like configuration rather than being selected freely throughout the search domain.

In this version, we added extra soft penalties to the geometry-aware acquisition function. The pairwise distances between contact points are encouraged to remain between $10\,\mathrm{mm}$ and $80\,\mathrm{mm}$ using a spacing-penalty weight of 20.0. Another term, with a weight of 2.0, encourages the four contacts to be roughly in a straight line. These terms do not constitute a full kinematic model of a robotic hand, but they offer a simple way to test the method under hand-like spatial constraints in the contact setup.

5.2.2. Surface-Based Extended Warm-Up Control

An additional experiment is performed for the compared method based on the combined GPIS–PFC formulation of de Farias et al. [1]. In their work, the exploration algorithm is initialized with prior samples generated before the exploration process begins. To make the comparison more balanced, 100 warm-up configurations are sampled directly from the object surface rather than the standard 50 uniformly random configurations. The purpose is to provide the surrogate model with more informative initial data before acquisition-guided optimization begins. All other experimental parameters remain as defined in Section 5.1.4 and Table 5.2. The geometry-aware acquisition function method retains its standard 50 random warm-up evaluations.

5.3. Evaluation Measures

This section defines the measures used to evaluate the experimental results. The quantitative analysis considers the GPIS surface cost, the estimated PFC, and the iterations required to reach selected PFC thresholds. Changes in the GPIS reconstruction are evaluated qualitatively.

5.3.1. GPIS Surface Cost

The GPIS surface cost, as defined in Chapter 4, is recorded for each joint contact configuration that is evaluated. Lower values mean the selected contacts are more consistent with the current GPIS surface estimate. When analyzing the search process, both the value at each iteration and the best value found so far are considered. The surface probability is calculated using a tolerance of $\varepsilon = 0.10$. This value indicates how well the selected contacts align with the current GPIS posterior.

5.3.2. Probability of Force Closure

The estimated Probability of Force Closure, $\widehat{P}_{\mathrm{FC}}(\mathbf{x}_t)$, is recorded for every evaluated joint contact configuration. Higher values indicate a larger fraction of Monte Carlo realizations satisfying force closure. Configurations for which the surface gate is not satisfied are assigned a PFC value of zero, as described in Section 5.1.4.

Since individual evaluations may explore configurations with lower grasp quality, the best PFC observed up to each iteration is also considered,

$$P_{\mathrm{FC},t}^{\mathrm{best}} = \max_{1 \le j \le t} \widehat{P}_{\mathrm{FC}}(\mathbf{x}_j). \tag{5.3}$$

This allows the analysis to distinguish between the PFC of the configuration evaluated at a particular iteration and the best grasp quality discovered up to that point.

5.3.3. Time to PFC Thresholds

To quantify how quickly each optimization run discovers grasp configurations with high estimated force-closure probability, the first iterations at which the PFC reaches 0.5 and 0.8 are recorded. These are defined as

$$t_{0.5}^{\mathrm{PFC}} = \min \left\{ t : \widehat{P}_{\mathrm{FC}}(\mathbf{x}_t) \geq 0.5 \right\}, \tag{5.4}$$

and

$$t_{0.8}^{\mathrm{PFC}} = \min \left\{ t : \widehat{P}_{\mathrm{FC}}(\mathbf{x}_t) \geq 0.8 \right\}. \tag{5.5}$$

The warm-up phase is not included when determining these first-hitting iterations. If a run does not reach a given PFC threshold within the 350 evaluations, that threshold is reported as not reached.

5.4. Experimental Results

This section presents the results of the Geometry-Aware Acquisition Function and the compared acquisition baselines. The Geometry-Aware Acquisition Function is analyzed first with respect to the exploration parameter, followed by the L-BFGS restart strategy. The baseline acquisition results are then reported separately. Finally, the hand-inspired contact experiment and the extended surface-based warm-up experiment are presented.

All five random seeds are included in the quantitative summaries. For each experimental configuration, $N_{0.5}$ and $N_{0.8}$ denote the number of runs that reach PFC values of at least 0.5 and 0.8, respectively. The reported first-hitting iterations are the medians over the runs that successfully reach the corresponding threshold. If none of the runs reaches a threshold, the first-hitting iteration is reported as "–". For all quantitative analyses in this chapter, the initial warm-up evaluations are excluded. Accordingly, PFC thresholds, best PFC values, and minimum GPIS surface costs are evaluated only after warm-up is finished. The original iteration numbering is retained throughout the analysis. All experimental configurations and parameter settings used in these evaluations are described in Section 5.2.

5.4.1. Geometry-Aware Acquisition Function Results

The Geometry-Aware Acquisition Function is evaluated using three exploration settings, $\kappa : 4.0 \rightarrow 0.8$, $\kappa = 2$, and $\kappa = 3$. All three settings are evaluated with the `last` restart

strategy on all five objects, resulting in 75 runs. The `best` restart strategy is additionally evaluated for the adaptive setting and fixed $\kappa = 2$ on the mustard bottle, tuna fish can, and cracker box, resulting in 30 additional runs.

The effect of the exploration parameter is evaluated first while keeping the restart strategy fixed to `last`. Results for the three exploration settings are summarized in Table 5.4.

Table 5.4.: Results of the Geometry-Aware Acquisition Function for the three exploration settings using the `last` restart strategy. Each row contains five random-seed runs. $N_{0.5}$ and $N_{0.8}$ represent the number of runs reaching the corresponding PFC threshold. Median first-hitting iterations are calculated only over successful runs. Median best PFC denotes the median of the maximum PFC values obtained in the five runs, while median minimum surface cost denotes the median of the lowest GPIS surface costs obtained in the five runs. Detailed run-level results are provided in Appendix A.1.1.

Object	κ	$\mathbf{N_{0.5}}$	**Median** $t_{0.5}^{\mathrm{PFC}} \downarrow$	$\mathbf{N_{0.8}}$	**Median** $t_{0.8}^{\mathrm{PFC}} \downarrow$	**Median best PFC** $\uparrow$	**Median min. surface cost** $\downarrow$
Mustard bottle	$4.0 \rightarrow 0.8$	5/5	55	5/5	79	1.0	0.000
Mustard bottle	2	5/5	55	5/5	**68**	1.0	0.000
Mustard bottle	3	5/5	61	5/5	79	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	5/5	61	5/5	127	1.0	0.000
Tuna fish can	2	5/5	65	5/5	143	1.0	0.000
Tuna fish can	3	5/5	61	5/5	**100**	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	5/5	56	5/5	**56**	1.0	0.000
Cracker box	2	5/5	60	5/5	60	1.0	0.000
Cracker box	3	5/5	74	5/5	75	1.0	0.000
Mini soccer ball	$4.0 \rightarrow 0.8$	5/5	53	5/5	65	1.0	0.000
Mini soccer ball	2	5/5	52	5/5	**54**	1.0	0.000
Mini soccer ball	3	5/5	54	5/5	57	1.0	0.000
Rubik's cube	$4.0 \rightarrow 0.8$	5/5	53	5/5	62	1.0	0.000
Rubik's cube	2	5/5	51	5/5	59	1.0	0.000
Rubik's cube	3	5/5	52	5/5	**58**	1.0	0.000

All three exploration settings reached both PFC thresholds in all five random-seed runs for every tested object. The median best PFC was 1.0, and the median minimum surface cost was zero for all object parameter combinations. The main differences between the exploration settings are therefore reflected in the number of iterations required to reach the PFC thresholds. For the 0.8 PFC threshold, fixed $\kappa = 2$ reached the lowest median first-hitting iteration on the mustard bottle and mini soccer ball, while fixed $\kappa = 3$ was fastest on the tuna fish can and Rubik's cube. The adaptive $\kappa : 4.0 \rightarrow 0.8$ setting reached the threshold earliest on the cracker box. Figures 5.2–5.4 illustrate the optimization behavior for the cracker box using seed 5 and the `last` restart strategy. This run is used as a representative example because its first-hitting iterations correspond to the median values obtained over the five random seeds for all three exploration settings.

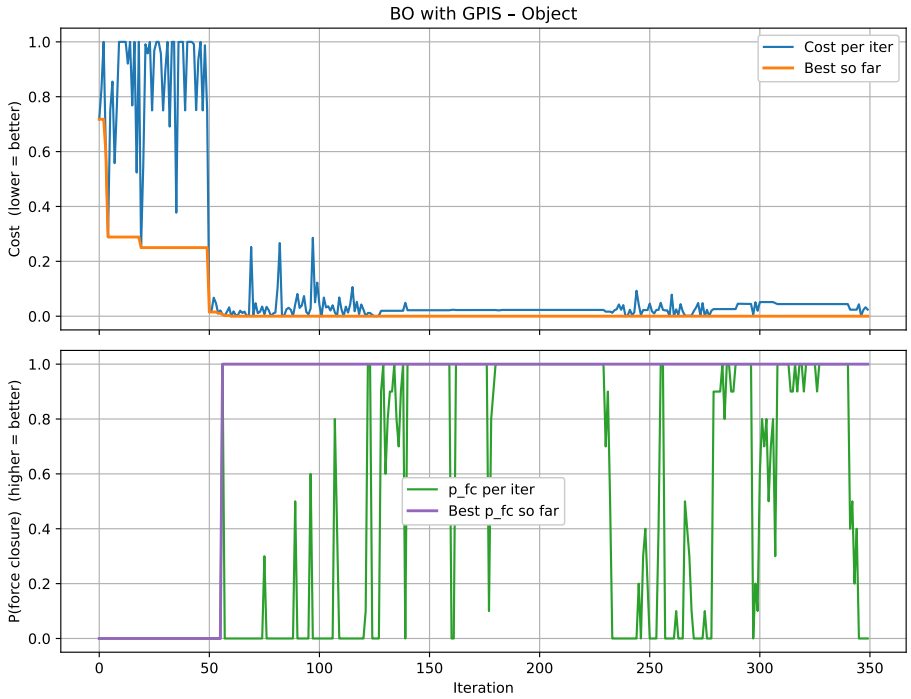


Figure 5.2.: Bayesian optimization results for the cracker box using the adaptive κ strategy, where κ decreases from 4.0 to 0.8, with the last-restart strategy. The upper plot shows the GPIS-based cost for each iteration together with the best cost found so far, while the lower plot shows the estimated Probability of Force Closure (PFC) for each iteration and the best PFC obtained so far.

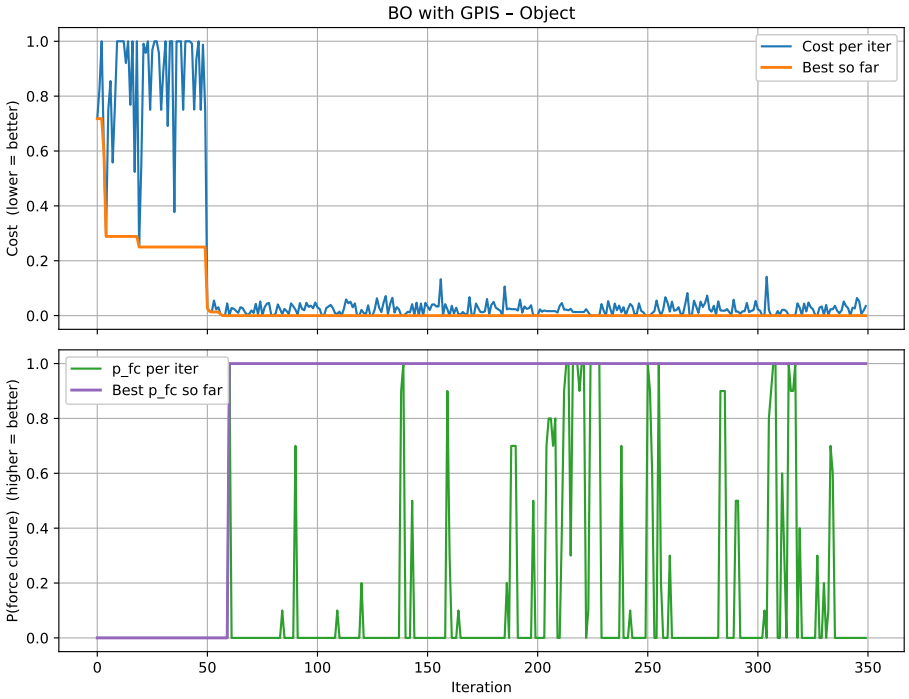


Figure 5.3.: Bayesian optimization results for the cracker box using a fixed $\kappa = 2$ with the last-restart strategy. The upper plot shows the GPIS-based cost for each iteration together with the best cost found so far, while the lower plot shows the estimated Probability of Force Closure (PFC) for each iteration and the best PFC obtained so far.

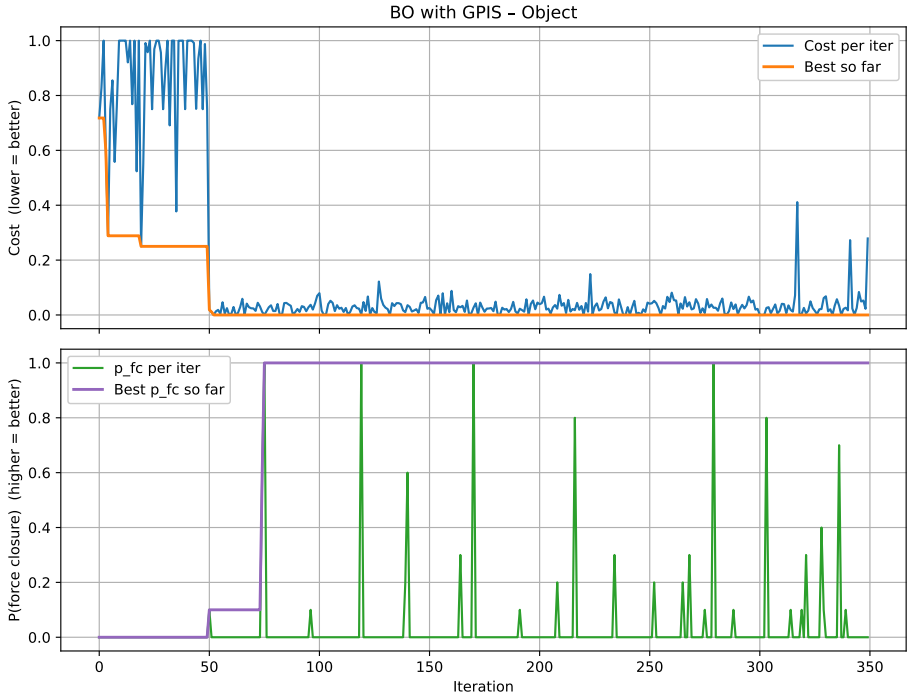


Figure 5.4.: Bayesian optimization results for the cracker box using a fixed $\kappa = 3$ with the last-restart strategy. The upper plot shows the GPIS-based cost for each iteration together with the best cost found so far, while the lower plot shows the estimated Probability of Force Closure (PFC) for each iteration and the best PFC obtained so far.

The influence of the L-BFGS restart strategy is evaluated by comparing `last` and `best` for the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting and fixed $\kappa = 2$. The results are summarized in Table 5.5. The comparison is performed separately for the mustard bottle, tuna fish can, and cracker box, for which both restart strategies were evaluated. For each object and exploration setting, the results are summarized over the same five random seeds.

Table 5.5.: Comparison of the `last` and `best` L-BFGS restart strategies for the adaptive $\kappa : 4.0 \rightarrow 0.8$ and fixed $\kappa = 2$ exploration settings. Each entry summarizes five random-seed runs for the same object and exploration setting. $N_{0.5}$ and $N_{0.8}$ denote the number of runs reaching PFC thresholds of 0.5 and 0.8, respectively. Values are reported as `last/best`. Detailed run-level results for the `last` and `best` restart strategies are provided in Appendix A.1.1 and Appendix A.1.2.

Object	κ	$N_{0.5}$	Median $t_{0.5}^{\mathrm{PFC}} \downarrow$	$N_{0.8}$	Median $t_{0.8}^{\mathrm{PFC}} \downarrow$	Median best PFC ↑	Median min. surface cost ↓
Mustard bottle	$4.0 \rightarrow 0.8$	5/5 / 5/5	55 / 55	5/5 / 5/5	79 / 79	1.0 / 1.0	0.000 / 0.000
Mustard bottle	2	5/5 / 5/5	55 / 58	5/5 / 5/5	**68** / 72	1.0 / 1.0	0.000 / 0.000
Tuna fish can	$4.0 \rightarrow 0.8$	5/5 / 5/5	61 / 62	5/5 / 5/5	127 / **111**	1.0 / 1.0	0.000 / 0.000
Tuna fish can	2	5/5 / 5/5	65 / 61	5/5 / 5/5	143 / **107**	1.0 / 1.0	0.000 / 0.000
Cracker box	$4.0 \rightarrow 0.8$	5/5 / 5/5	56 / 56	5/5 / 5/5	56 / 56	1.0 / 1.0	0.000 / 0.000
Cracker box	2	5/5 / 5/5	60 / 60	5/5 / 5/5	60 / 60	1.0 / 1.0	0.000 / 0.000

Both restart strategies reached both PFC thresholds in all five runs for every evaluated object–parameter combination. The median best PFC was 1.0, and the median minimum surface cost was zero throughout the comparison. For the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting, the two restart strategies produced identical median first-hitting iterations on the mustard bottle and cracker box, while `best` reached the 0.8 threshold earlier on the tuna fish can, with median iterations of 111 and 127 for `best` and `last`. For fixed $\kappa = 2$, `last` was slightly earlier on the mustard bottle, `best` was earlier on the tuna fish can, and the two strategies produced identical medians on the cracker box.

The qualitative evolution of the GPIS prediction is illustrated in Figures 5.5 and 5.6 using the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting with the `last` restart strategy. For each object, the left image shows the GPIS prediction before Bayesian optimization, while the right image shows the prediction after 350 evaluations. The same viewpoint is used before and after optimization to allow a direct visual comparison. The colored markers visible in the final predictions indicate contact points selected for force-closure evaluation during optimization. Since each object is shown only from a single viewpoint, only the contact points visible from that side of the object are displayed.

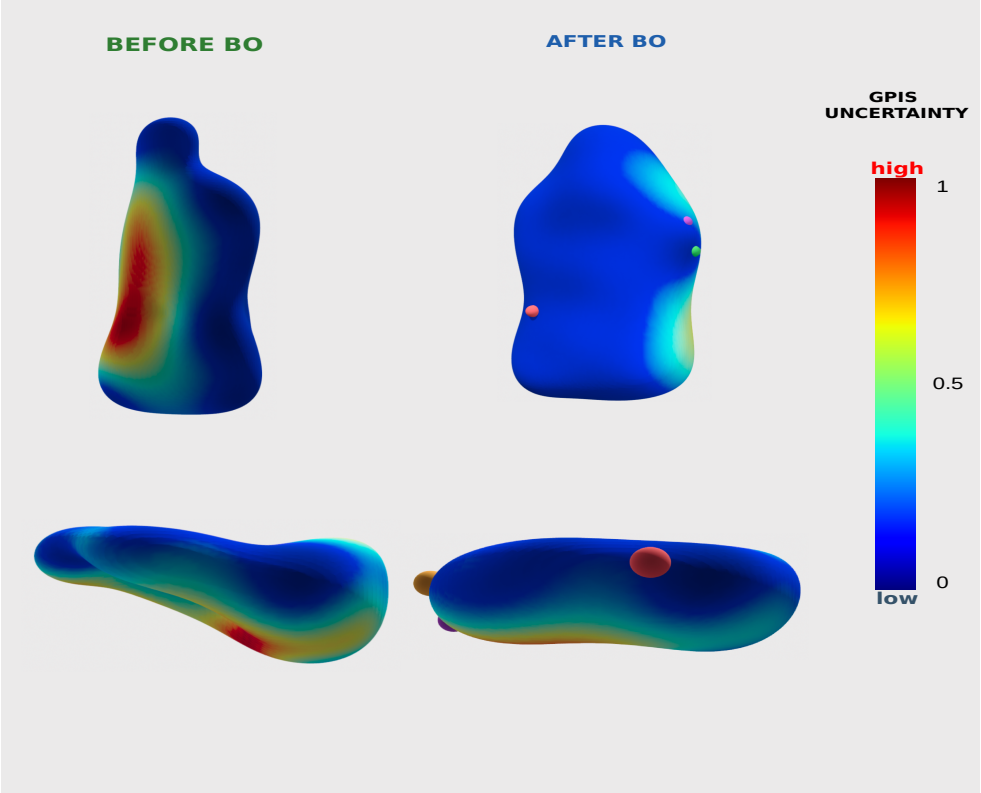


Figure 5.5.: GPIS predictions before and after Bayesian optimization for the mustard bottle (top) and tuna fish can (bottom). For each object, the initial prediction is shown on the left, and the prediction after 350 evaluations is shown on the right. The surface color represents the normalized GPIS predictive variance, with blue indicating lower uncertainty and red indicating higher uncertainty.

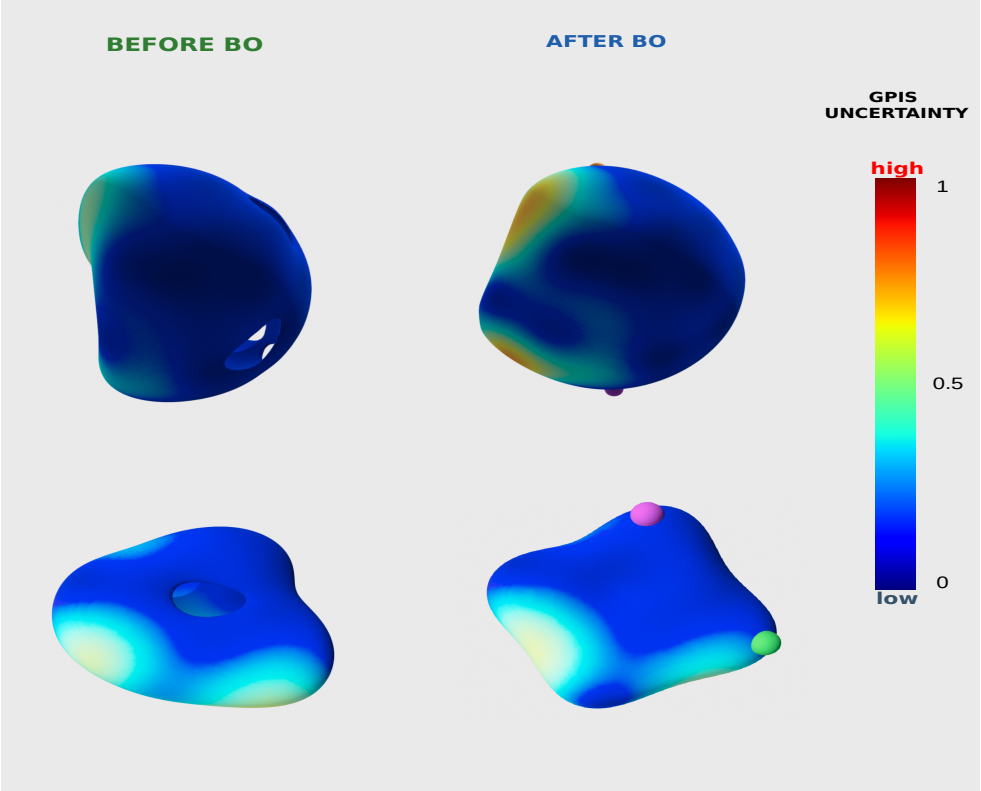


Figure 5.6.: GPIS predictions before and after Bayesian optimization for the mini soccer ball (top) and Rubik's cube (bottom). For each object, the initial prediction is shown on the left, and the prediction after 350 evaluations is shown on the right. The surface color represents the normalized GPIS predictive variance, with blue indicating lower uncertainty and red indicating higher uncertainty.

5.4.2. Baseline Acquisition Strategy Results

This subsection presents the results for the three baseline acquisition strategies defined in Section 5.2: PFC-only UCB, Expected Improvement, and the combined-target baseline based on de Farias et al. [1].

Results for the PFC-only UCB and Expected Improvement baselines are summarized in Tables 5.6 and 5.7, respectively. The parameter settings for both baselines are defined in Section 5.2. Each setting is evaluated on all five test objects using five random seeds, resulting in 75 runs for each baseline.

Table 5.6.: Results of the PFC-only UCB acquisition for the three exploration settings. Each row summarizes five random-seed runs.

Object	κ	$N_{0.5}$	Median $t_{0.5}^{PFC}\downarrow$	$N_{0.8}$	Median $t_{0.8}^{PFC}\downarrow$	Median best PFC $\uparrow$	Median min. surface cost $\downarrow$
Mustard bottle	0	0/5	– –	0/5	– –	0.0	**0.517**
Mustard bottle	2	0/5	– –	0/5	– –	0.0	1.0
Mustard bottle	3	0/5	– –	0/5	– –	0.0	1.0
Tuna fish can	0	0/5	– –	0/5	– –	0.0	**0.500**
Tuna fish can	2	0/5	– –	0/5	– –	0.0	1.0
Tuna fish can	3	0/5	– –	0/5	– –	0.0	1.0
Cracker box	0	0/5	– –	0/5	– –	0.0	**0.500**
Cracker box	2	0/5	– –	0/5	– –	0.0	1.0
Cracker box	3	0/5	– –	0/5	– –	0.0	1.0
Mini soccer ball	0	0/5	– –	0/5	– –	0.0	**0.510**
Mini soccer ball	2	0/5	– –	0/5	– –	0.0	1.0
Mini soccer ball	3	0/5	– –	0/5	– –	0.0	1.0
Rubik's cube	0	0/5	– –	0/5	– –	0.0	**0.292**
Rubik's cube	2	0/5	– –	0/5	– –	0.0	1.0
Rubik's cube	3	0/5	– –	0/5	– –	0.0	1.0

Table 5.7.: Results of the Expected Improvement acquisition for the three values of ξ. Each row summarizes five random-seed runs.

Object	$\boldsymbol{\xi}$	$\mathbf{N_{0.5}}$	**Median** $t_{0.5}^{\mathrm{PFC}} \downarrow$	$\mathbf{N_{0.8}}$	**Median** $t_{0.8}^{\mathrm{PFC}} \downarrow$	**Median best PFC** $\uparrow$	**Median min. surface cost** $\downarrow$
Mustard bottle	0.0	0/5	−−	0/5	−−	0.0	1.000
Mustard bottle	0.1	0/5	−−	0/5	−−	0.0	1.000
Mustard bottle	0.2	0/5	−−	0/5	−−	0.0	1.000
Tuna fish can	0.0	0/5	−−	0/5	−−	0.0	1.000
Tuna fish can	0.1	0/5	−−	0/5	−−	0.0	1.000
Tuna fish can	0.2	0/5	−−	0/5	−−	0.0	1.000
Cracker box	0.0	0/5	−−	0/5	−−	0.0	1.000
Cracker box	0.1	0/5	−−	0/5	−−	0.0	1.000
Cracker box	0.2	0/5	−−	0/5	−−	0.0	1.000
Mini soccer ball	0.0	0/5	−−	0/5	−−	0.0	1.000
Mini soccer ball	0.1	0/5	−−	0/5	−−	0.0	1.000
Mini soccer ball	0.2	0/5	−−	0/5	−−	0.0	1.000
Rubik's cube	0.0	0/5	−−	0/5	−−	0.0	1.000
Rubik's cube	0.1	0/5	−−	0/5	−−	0.0	1.000
Rubik's cube	0.2	0/5	−−	0/5	−−	0.0	1.000

Results for combined-target baseline based on de Farias et al. [1] are summarized in Table 5.8. The two parameter settings defined in Section 5.2 are evaluated on the mustard bottle, tuna fish can, and cracker box using five random seeds, resulting in 30 runs in total.

Table 5.8.: Results of the combined-target baseline based on de Farias et al. [1]. Each row summarizes five random-seed runs. Detailed run-level results are provided in Appendix A.2.

Object	κ	$N_{0.5}$	Median $t_{0.5}^{PFC}$ ↓	$N_{0.8}$	Median $t_{0.8}^{PFC}$ ↓	Median best PFC ↑	Median min. surface cost ↓
Mustard bottle	$4.0 \rightarrow 0.8$	$0/5$	$--$	$0/5$	$--$	0.0	1.000
Mustard bottle	2	$0/5$	$--$	$0/5$	$--$	0.0	**0.750**
Tuna fish can	$4.0 \rightarrow 0.8$	$0/5$	$--$	$0/5$	$--$	0.0	1.000
Tuna fish can	2	$0/5$	$--$	$0/5$	$--$	0.0	**0.500**
Cracker box	$4.0 \rightarrow 0.8$	$0/5$	$--$	$0/5$	$--$	0.0	1.000
Cracker box	2	$0/5$	$--$	$0/5$	$--$	0.0	**0.509**

None of the baseline runs reached either PFC threshold. For PFC-only UCB, $\kappa = 0$ was the only setting where we saw a reduced median minimum surface cost below 1.0 for all five objects. Expected Improvement produced a median best PFC of 0.0 and a median minimum surface cost of 1.0 for every tested object and parameter setting; therefore, there was nothing improved. For the combined-target baseline based on de Farias et al. [1], fixed $\kappa = 2$ reduced the median minimum surface cost on all three tested objects, but did not reach any force closure.

5.4.3. Hand-Inspired Contact Configuration Results

The hand-inspired configuration described in Section 5.2.1 is evaluated on the mustard bottle, tuna fish can, and cracker box using the adaptive $\kappa : 4.0 \rightarrow 0.8$ exploration setting and the `last` restart strategy. Five random seeds are used for each object, resulting in 15 runs in total. The results are summarized in Table 5.9.

Table 5.9.: Results of the hand-inspired contact experiment using the adaptive $\kappa : 4.0 \rightarrow 0.8$ exploration setting and the `last` restart strategy. Each row summarizes five random-seed runs. Detailed run-level results are provided in Appendix A.3.

Object	$N_{0.5}$	Median $t_{0.5}^{PFC} \downarrow$	$N_{0.8}$	Median $t_{0.8}^{PFC} \downarrow$	Median best PFC $\uparrow$	Median min. surface cost $\downarrow$
Mustard bottle	5/5	74	5/5	106	1.0	0.000
Tuna fish can	5/5	74	4/5	128	0.9	0.000
Cracker box	5/5	60	5/5	**72**	1.0	0.000

All five runs reached the 0.5 PFC threshold for each of the three objects. The 0.8 threshold was reached in all five runs for the mustard bottle and cracker box, and in four of the five runs for the tuna fish can. The median best PFC was 1.0 for the mustard bottle and cracker box and 0.9 for the tuna fish can, while the median minimum surface cost was zero for all three objects. A representative configuration for the mustard bottle is shown in Figure 5.7. One contact is located on one side of the object, while the remaining three contacts are grouped on the opposite side, resulting in a simplified hand-like contact arrangement.

5.4.4. Extended Surface-Based Warm-Up Results

The extended surface-based warm-up experiment described in Section 5.2.2 is evaluated on the mustard bottle and tuna fish can. The two exploration settings are tested using five random seeds and the `last` restart strategy. The results are summarized in Table 5.10.

Table 5.10.: Results of the extended surface-based warm-up experiment for the combined-target baseline based on de Farias et al. [1], using 100 surface-based warm-up evaluations. Each row summarizes five random-seed runs. Detailed run-level results are provided in Appendix A.4.

Object	κ	$N_{0.5}$	Median $t_{0.5}^{PFC} \downarrow$	$N_{0.8}$	Median $t_{0.8}^{PFC} \downarrow$	Median best PFC $\uparrow$	Median min. surface cost $\downarrow$
Mustard bottle	$4.0 \rightarrow 0.8$	0/5	$--$	0/5	$--$	0.0	0.250
Mustard bottle	2	0/5	$--$	0/5	$--$	0.0	**0.217**
Tuna fish can	$4.0 \rightarrow 0.8$	5/5	273	4/5	330	**0.8**	0.000
Tuna fish can	2	5/5	144	2/5	**258.5**	0.7	0.000

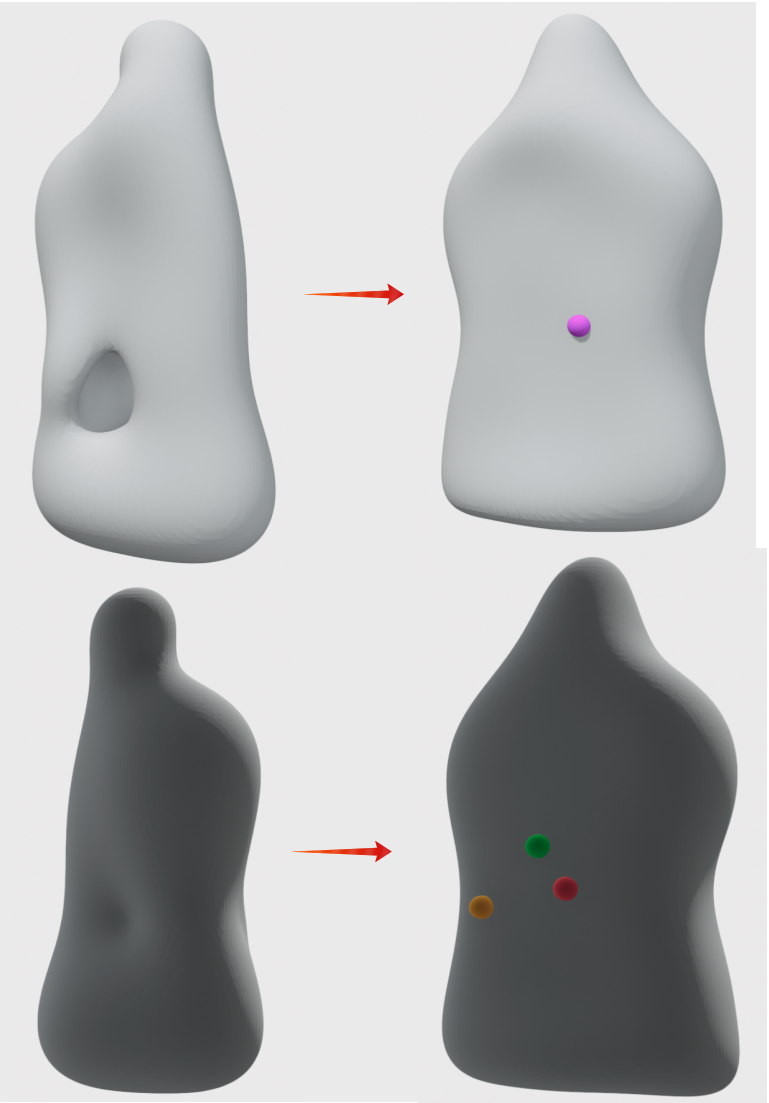


Figure 5.7.: Representative hand-inspired contact configuration on the mustard bottle. On the left side, we have the initial GPIS prediction, and the arrows indicate the transition after Bayesian Optimization, the updated GPIS model, and the reconstructed object surface with the selected contact configuration. The upper row shows the side containing the single opposing contact, while the lower row shows the opposite side containing the three grouped contacts.

6. Discussion

The experimental evaluation in Chapter 5 concludes experiments to see if direct geometric information from the evolving GPIS model can improve Bayesian optimization of multi-contact configurations on partially observed objects. The results show a clear difference between the Geometry-Aware Acquisition Function and the alternative acquisition strategies. The Geometry-Aware Acquisition Function reached both considered PFC thresholds consistently across the tested objects, whereas methods using only a surrogate model in the acquisition function, PFC-only UCB, Expected Improvement, and the method based on the combined GPIS–PFC formulation of de Farias et al. did not reach either threshold in the corresponding standard experiments. In this chapter, we discuss why such a difference occurs.

The contact-search problem considered in this thesis is particularly difficult because four three-dimensional contact positions are optimized independently, resulting in a 12-dimensional Cartesian search space. Furthermore, a useful configuration must satisfy two requirements simultaneously: the contacts must first correspond sufficiently well to the estimated object surface, and they must then form a configuration with a high probability of force closure. The following sections discuss how our acquisition function addresses these requirements, why the alternative acquisition strategies behave differently, and how these observations relate to robotic grasp planning.

6.1. Role of Direct Geometric Guidance

The central experimental result is the consistency of the Geometry-Aware Acquisition Function. With the `last` restart strategy, all 75 runs across the five objects and three exploration settings reached both $P_{\mathrm{FC}} \geq 0.5$ and $P_{\mathrm{FC}} \geq 0.8$. In addition, the median best PFC was 1.0 and the median minimum GPIS surface cost was zero for every object. Median values are calculated across all runs, which are reported in Table 5.4. This

behavior is particularly relevant considering the structure of the search space. Each of the four contacts can initially be placed independently anywhere inside the three-dimensional search domain. A geometrically useful configuration, therefore, requires four independent points to be placed simultaneously close to the estimated object surface. These configurations are found only in a limited part of the full 12-dimensional Cartesian search space. If we require these contacts to create a force-closure grasp, the number of useful configurations becomes even smaller.

The PFC surface gate makes this problem more pronounced. PFC is evaluated only when the GPIS surface cost satisfies $s_{\text{surf}} \leq 0.1$. Otherwise, the configuration receives a PFC value of zero. Therefore, before the optimization begins to produce geometrically valid contact configurations, many different points in the 12-dimensional space result in the same PFC observation. The PFC surrogate alone cannot distinguish whether such a zero value was produced because the contacts were close to the object but formed a poor grasp, or because one or more contacts were far from the estimated surface.

The Geometry-Aware Acquisition Function avoids depending exclusively on this initially weak PFC signal. The GPIS posterior is queried directly during candidate evaluation. Its predictive mean provides information about if the proposed point is expected to lie on the surface, while its predictive uncertainty provides information about regions whose geometry remains uncertain. This gives the acquisition function a meaningful direction even when the PFC surrogate has not yet observed useful force-closure values. This distinction is important because the PFC Gaussian-process surrogate in the present implementation is not supplied with prior PFC training data. It is initially empty and receives observations through the random warm-up and subsequent optimization evaluations. Nevertheless, after the standard 50 random warm-up evaluations, the Geometry-Aware Acquisition Function is able to use the GPIS component to guide the search towards the object and subsequently obtain informative PFC observations.

The two components, UCB from GPIS model and from the surrogate GP, therefore have complementary functions. The GPIS term addresses the geometric part of the problem by identifying points that lie on the surface, while the PFC term distinguishes among geometrically plausible configurations based on grasp stability. Once the points of the proposed configuration are on the surface and the surface gate is satisfied, the PFC surrogate becomes increasingly useful, and the search changes from primarily finding the object surface to simultaneously refining surface agreement and grasp quality.

Another useful property of this formulation is that changes to the GPIS model immediately affect acquisition values throughout the contact space. Whenever new geometric observations update the GPIS posterior, the GPIS acquisition component is recomputed

directly from the new posterior for every candidate. The optimization therefore has direct access to the latest geometric model without requiring a separate surrogate to relearn the full geometric effect from scalar-objective observations. However, we have to say that the GPIS model is not exactly a representation of the unknown object, at least at the beginning, since it has only part of the object. By the end of the optimization process, GPIS is able to represent almost full object geometries.

6.2. Comparison with Baseline Acquisition Strategies

The alternative acquisition experiments provide additional evidence that explicit geometric guidance is important in the present search problem. PFC-only UCB and Expected Improvement remove direct access to the GPIS posterior during candidate selection, while the method based on de Farias et al. includes geometric information indirectly through a combined scalar target. However, in all three approaches, only a surrogate model is used to guide the acquisition function. These differences result in substantially different optimization behavior.

6.2.1. PFC-Only UCB

During the 75 PFC-only UCB runs, none of them reached either of the considered PFC thresholds. This result cannot be explained simply by the choice of UCB, because we used UCB also in our Geometry-Aware Acquisition Function, both for the PFC surrogate model and for the GPIS model. The important difference is the information available to the acquisition function, and in this case there is no information about the object in the acquisition function. In PFC-only UCB, candidate selection is based exclusively on the predictive mean and uncertainty of the PFC surrogate. At the beginning of the search, however, the surrogate is populated largely by configurations with a PFC of zero because the surface gate is not satisfied. As a result, the surrogate contains little information describing how the individual contacts should move towards the object surface.

In the present problem, uncertainty of the PFC surrogate is not the same as uncertainty about the object geometry. A configuration can have a high UCB value because it lies far from previously evaluated configurations while still placing its contacts in geometrically irrelevant regions of the search domain. The optimization therefore explores the 12-dimensional PFC search space without any direct information about where valid surface

contacts are located. This behavior is consistent with experiments with fixed $\kappa = 2$ and $\kappa = 3$, in which the optimization did not yield meaningful surface-cost improvements. The uncertainty term encourages exploration, but it provides no information about which unexplored regions contain the object surface. This difficulty is further increased by the surface gate: if the selected contact points are not sufficiently close to the GPIS surface, the PFC value is set directly to zero. This makes the PFC surrogate even less informative during the early stages of optimization, since many geometrically invalid configurations produce the same zero-valued observation. However, evaluating force closure for contacts that do not lie on the object surface would not be physically meaningful, so this gating step is necessary despite the additional difficulty it introduces for the surrogate model.

On the other hand, using $\kappa = 0$, we obtained some results. Some reduction in GPIS surface cost was observed in this setting, particularly for the Rubik's cube, but the PFC remained zero. Since the acquisition contains no GPIS term, these surface improvements do not represent explicit surface-directed search. Configurations selected through the PFC surrogate can incidentally produce useful geometric observations, which happened in this case, resulting in an improved GPIS model. This was not sufficient to produce a complete four-contact configuration that satisfied the surface gate and provided a useful force-closure signal.

The comparison with our Geometry-Aware Acquisition Function therefore suggests that increasing generic Bayesian-optimization exploration is not sufficient for this problem. The search benefits specifically from exploration that is informed by the geometry of the partially observed object.

6.2.2. Expected Improvement

Expected Improvement exhibited an even more uniform failure mode. For all 75 evaluated runs and all three values of ξ, the PFC remained zero and the GPIS surface cost remained at 1.0. EI selects candidates according to their expected improvement over the best objective value observed so far. In the present experiments, however, the PFC surface gate causes many observations to be zero. Therefore, there is only a small number of signals from which the acquisition can provide a direction towards useful contact configurations.

Predictive uncertainty remains part of EI, but, as with PFC-only UCB, uncertainty in the PFC surrogate does not indicate the location of the object's surface. The resulting search therefore has no explicit mechanism for moving four independent contact positions towards geometrically valid locations.

The failure of Expected Improvement in these experiments should not be interpreted as evidence that EI is generally unsuitable for robotic grasping or Bayesian optimization. The difficulty comes from using EI only on the PFC surrogate in this particular problem. Before EI can receive a meaningful PFC signal, the selected contact points must first satisfy the GPIS surface gate. Otherwise, the PFC value is set to zero. The Geometry-Aware Acquisition Function avoids this problem by using the GPIS posterior directly in the acquisition function.

6.2.3. Method Based on the Combined GPIS–PFC Formulation of de Farias et al.

The comparison with the method based on de Farias et al. is more informative than the PFC-only comparisons because geometric information is present and incorporated in the optimization target. Following their general formulation, a GPIS-derived surface term and PFC are combined into a single scalar observation, which is then modeled using one Gaussian-process surrogate [1]. In the implementation used in this thesis, UCB is applied to this surrogate, whereas the original work of de Farias et al. uses Expected Improvement.

Despite the inclusion of geometric information, none of the 30 standard runs of this compared method reached either PFC threshold. For the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting, the reported median surface cost remained at 1.0. Fixed $\kappa = 2$ produced some geometric improvement, with median minimum surface costs of 0.750, 0.500, and 0.509 for the mustard bottle, tuna fish can, and cracker box, respectively, but PFC remained zero. These results illustrate an important distinction between including GPIS information in the optimization target and using the GPIS posterior directly inside the acquisition function. In the combined-target approach, geometry and grasp quality are first compressed into one scalar value. The BO surrogate must then learn how this scalar target varies over the complete 12-dimensional configuration space.

In contrast, the Geometry-Aware Acquisition Function does not require the PFC surrogate to learn the object's surface location because the existing GPIS model already contains this information. The acquisition directly evaluates the posterior mean and uncertainty at each candidate contact point and combines this geometric score with the separately modeled PFC signal. This distinction becomes especially important when the GPIS is updated during exploration. In the Geometry-Aware Acquisition Function, a new geometric observation changes the GPIS posterior and therefore immediately changes the GPIS acquisition value, which is used for the next iteration in choosing new candidates. In a single-surrogate combined-target formulation, the BO surrogate model does not get the complete

updated GPIS posterior. Its knowledge of the combined objective is obtained through the scalar targets observed at previously evaluated configurations. It must therefore learn a more complex mapping from a 12-dimensional grasp configuration to a quantity that incorporates both geometric and force-closure effects.

The results suggest that keeping the GPIS and PFC information separate and using both directly in the acquisition function is beneficial for the contact-search problem considered in this thesis. However, we have to say that our proposed method and the experiments were only done on simulation, while the method proposed by de Farias et al. was developed and evaluated in a robotic setting with additional kinematic and physical constraints.

6.3. Difference from the Robotic Setting of de Farias et al.

The work in the paper of de Farias et al. optimizes grasp configurations for a three-finger Schunk SDH robotic hand [1]. Although their optimization space is also reported as twelve-dimensional, it is not equivalent to the twelve independent coordinates used in our present experiments. In this thesis, each of the four contact points can be placed independently at any location inside the three-dimensional bounding box. Before hand-inspired penalties are introduced, there is no kinematic relationship between the contacts.

A physical robotic hand naturally removes many such configurations. In the work of de Farias et al., the positions of the fingers and wrist are restricted by the physical hand. With these constraints, the optimization problem, i.e., the search space, is actually changed. Once one finger of a real hand is placed on an object, the remaining fingers cannot move independently to arbitrary locations elsewhere in the search domain. Hand kinematics therefore provide a strong structural prior about which combinations of contact positions are physically meaningful. Without a physical hand model, the standard experiments use an unconstrained search in which the contact points can move independently. With an unconstrained search, it is possible to evaluate the acquisition strategy itself, but it also produces a much larger set of geometrically and physically irrelevant configurations. The fact that the Geometry-Aware Acquisition Function finds four surface contacts and, at the same time, high-PFC configurations under this parameterization is therefore an important result. It does so without relying on a complete hand model to restrict the search to plausible grasps. This aspect is important when jointly reconstructing the object and searching for force-closure contacts in an unconstrained space. In a real robotic system, however, the hand kinematics already restrict the set of feasible contact configurations, so this particular advantage would likely be less significant.

Initialization is another relevant difference. De Farias et al. [1] initialize their exploration procedure using prior samples within a robotic grasping setup, where the hand kinematics already restrict the sampled configurations to a structured set of physically meaningful grasps. In the standard experiments of this thesis, the PFC surrogate starts without prior PFC observations and is initialized only through 50 uniformly sampled joint contact configurations in the unconstrained 12-dimensional Cartesian search space.

This makes the combined-surrogate approach of De Farias et al. in our setup harder to optimize, since it has to learn a single GPIS-PFC objective over a large and only weakly constrained contact-search space, without the same prior information or kinematic restrictions provided by a robotic hand as it was in their setup.

6.4. Effect of Surface-Based Initialization

The extended surface-based warm-up experiment was used to test whether more informative initial geometric data could improve the combined-surrogate method proposed by de Fearias et al. Instead of 50 random warm-up configurations, the method was provided with 100 surface-based evaluations before acquisition-guided optimization.

The effect was strongly object-dependent. On the mustard bottle, the median minimum surface cost improved from 1.000 to 0.250 for the adaptive setting and from 0.750 to 0.217 for fixed $\kappa = 2$, but none of the runs reached $P_{\mathrm{FC}} \geq 0.5$. This shows that improving surface information alone does not necessarily lead to a successful force-closure configuration. In the experiment with the tuna fish, we can see a different behavior. With the adaptive setting, all five runs reached $P_{\mathrm{FC}} \geq 0.5$ and four reached $P_{\mathrm{FC}} \geq 0.8$. For fixed $\kappa = 2$, all five reached the lower threshold and two reached the higher threshold. This is a clear improvement over the standard de Farias-based experiments that we ran, where none of the runs reached either threshold.

However, the extended warm-up still did not match our Geometry-Aware Acquisition Function. For the adaptive setting, the median $t_{0.8}^{\mathrm{PFC}}$ was 330, compared with 127 for the proposed method. For fixed $\kappa = 2$, the corresponding values were 258.5 and 143. The results therefore show that better initialization can help the combined surrogate, but its effect is object-dependent and less consistent than directly using the evolving GPIS posterior in the acquisition function.

The extended warm-up experiment, in the case of De Farias et al. approach, therefore supports two conclusions. First, the combined-surrogate formulation is sensitive to the

quality of its initial observations. Second, providing additional geometric information can improve its behavior, but this does not provide the same consistency as directly using the evolving GPIS posterior during acquisition.

We tried the same extended surface-based warm-up with the PFC-only UCB and Expected Improvement baselines. This initialization did not produce a meaningful improvement for either PFC-only UCB or Expected Improvement, and these additional trials are therefore not analyzed separately. The Geometry-Aware Acquisition Function achieves successful search using only the standard 50 random warm-up samples and does not require surface-based initial contact configurations.

6.5. Sensitivity of the Geometry-Aware Acquisition Function

The Geometry-Aware Acquisition Function was also evaluated with different exploration parameters and L-BFGS restart strategies. In contrast to the large difference between acquisition strategies, these internal design choices mainly affected the speed of convergence rather than whether a successful configuration was eventually found.

6.5.1. Exploration Parameter

All three tested exploration settings reached both PFC thresholds in every standard run. The median best PFC and minimum surface cost were also identical across all object–parameter combinations. The differences were therefore primarily visible in the first-hitting iterations.

No single value of κ was consistently fastest. Fixed $\kappa = 2$ reached the 0.8 threshold earlier on the mustard bottle and mini soccer ball, fixed $\kappa = 3$ was faster on the tuna fish can and Rubik's cube, and the adaptive setting was fastest on the cracker box. This indicates that the exact exploration weight affects how efficiently the method moves through the search space, but the direct geometric component makes the method relatively robust to the tested values of κ. The adaptive schedule should therefore not be interpreted as uniformly superior. Its advantage is that it provides a mechanism for changing the relative importance of uncertainty as optimization progresses, whereas the fixed settings use the same exploration weight throughout. In the present experiments, however, all three settings were sufficiently effective to solve the tested contact-search problems.

6.5.2. Restart Strategy

The L-BFGS restart comparison produced a similar result. Both `last` and `best` reached both PFC thresholds in every evaluated run, and both produced a median best PFC of 1.0 and a minimum surface cost of zero.

The main differences again occurred in convergence speed. The most noticeable effect was observed for the tuna fish can, where the `best` restart reduced the median $t_{0.8}^{\mathrm{PFC}}$ from 127 to 111 for the adaptive setting and from 143 to 107 for fixed $\kappa = 2$. On the cracker box, the median first-hitting iterations were identical for the two restart strategies, while on the mustard bottle the differences were small.

The restart results therefore do not establish one strategy as generally superior. They instead show that retaining the best previous configuration can occasionally accelerate local optimization, particularly in a more difficult search landscape, but that our acquisition function remains successful with either strategy.

6.6. Hand-Inspired Contact Search

The standard experiments allow the four contact points to move independently. The hand-inspired experiment tests whether the Geometry-Aware Acquisition Function remains effective also when additional constraints are introduced between the contacts.

All 15 hand-inspired runs reached $P_{\mathrm{FC}} \geq 0.5$, and 14 of 15 reached $P_{\mathrm{FC}} \geq 0.8$. The median minimum surface cost was zero for all three tested objects, showing that the additional spacing and collinearity penalties did not prevent the method from finding high-quality configurations. However, the restrictions increased the search time in some cases. For the mustard bottle, the median $t_{0.8}^{\mathrm{PFC}}$ increased from 79 to 106, and for the cracker box from 56 to 72. The tuna fish can showed almost no change, with median values of 127 and 128 among successful runs.

This experiment provides a first step towards a more realistic robotic setting in which contact positions cannot move independently. However, the hand-inspired penalties are only a simplified approximation of a real hand and do not include joint limits, complete hand kinematics, and collision constraints.

7. Conclusion

This thesis addressed multi-contact grasp search on partially observed objects using Bayesian optimization. Our approach and its Geometry-Aware Acquisition Function combine geometric information from an evolving Gaussian Process Implicit Surface (GPIS) with a probabilistic force-closure GP model, enabling surface exploration and grasp search to be performed within a single optimization process. This chapter summarizes the main findings of the work and outlines directions for future development.

7.1. Main Findings

The experimental results showed that the Geometry-Aware Acquisition Function was able to consistently find contact configurations with a high probability of force closure across the tested objects. In contrast, the compared PFC-only approaches were not able to provide sufficient guidance in the large contact search space, particularly before valid surface contacts had been identified.

The experiments based on the combined GPIS–PFC formulation of de Farias et al. showed that geometric information can also be incorporated through a combined surrogate, but its performance was more dependent on the quality of the initial data. The surface-based warm-up experiments confirmed that better initial geometric information can improve this behavior, although the improvement was not consistent across objects.

Overall, the results indicate that using GPIS information directly during acquisition is beneficial when surface exploration and force-closure search are performed together. The hand-inspired experiments further showed that this approach remains effective when additional restrictions are introduced between the contact points, providing a first step towards a more realistic robotic setting.

7.2. Limitations and Future Work

The main limitation of the current evaluation is that all experiments were performed in a simplified simulation environment. In this situation, it allows us to compare without possible failures caused by sensing, calibration, robot control, or contact dynamics, but it does not demonstrate how the Geometry-Aware Acquisition Function would behave in a robotic simulation or on a physical robotic system. Therefore, the next step is to evaluate the method in a simulation with a robot before transferring it to a real robotic hand. In this way, using first the simulation robot, we introduce contact dynamics, friction, collision handling and robot kinematics in a controlled setting. Instead of optimizing four independent contact positions, the search could be performed over reachable hand poses or joint configurations.

For the real-world implementation, we would need some more reliable estimation of the contact information that is directly available in the current simulation. In particular, contact positions and surface normals would have to be obtained from tactile sensing, force-torque measurements, or their combination. This would directly affect both the GPIS update but also the PFC surogate model. Collision checking and feasible approach trajectories would also have to be considered before a candidate grasp could be executed on the physical system.

Another limitation is that the standard formulation in this thesis optimizes four three-dimensional contact positions independently and does not enforce the kinematics of a specific robotic hand. The hand-inspired experiment introduces one step forward , but it remains only a simplified approximation of physically reachable finger contacts. Future work could incorporate kinematic constraints directly into the search by optimizing over reachable hand poses or joint configurations instead of independent Cartesian contact positions. On the other hand, infeasible configurations should be rejected, and in some way, joint limits, self-collisions, and collisions with the environment should be checked and penalized.

For the problem considered in this thesis, it was also assumed that the object remains static throughout the optimization procedure. This allows observations collected at different iterations to be accumulated directly in a single GPIS model. We did not take into account If the object moves or is disturbed during exploration, and what effects it would have for the GPIS where we reconstruct our object, but also for the surrogate model which captures configurations with high PFC. Extending the method to non-static objects would

therefore require real-time object tracking together with corresponding updates of the GPIS model and the contact-search domain.

Further work could also investigate the surface gate used before PFC evaluation. The current hard threshold is physically meaningful because force closure should not be evaluated for contacts that do not lie sufficiently close to the estimated object surface, but it also produces many zero-valued PFC observations during the early search. A softer or uncertainty-dependent surface condition could therefore be considered to reduce this effect while still preventing clearly invalid contact configurations from contributing force-closure values.

Finally, the GPIS and PFC surrogate models could be studied in more detail. Due to time constraints, only the Matérn-$5/2$ kernel was used in the present experiments. Although this provides a consistent kernel choice throughout the evaluation, alternative kernels could be compared to determine their influence on object reconstruction and optimization behavior. In addition, a quantitative reconstruction metric could be introduced to evaluate the accuracy of the complete reconstructed object rather than only the agreement between selected contacts and the current GPIS surface. This would make it possible to analyze object reconstruction quality and grasp-search performance as separate outcomes.

Bibliography

[1] C. De Farias, N. Marturi, R. Stolkin, and Y. Bekiroglu, "Simultaneous tactile exploration and grasp refinement for unknown objects," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3349–3356, 2021.

[2] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, vol. 2. MIT press Cambridge, MA, 2006.

[3] C. Williams and C. Rasmussen, "Gaussian processes for regression," *Advances in neural information processing systems*, vol. 8, 1995.

[4] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.

[5] E. Schulz, M. Speekenbrink, and A. Krause, "A tutorial on gaussian process regression: Modelling, exploring, and exploiting functions," *Journal of mathematical psychology*, vol. 85, pp. 1–16, 2018.

[6] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global optimization*, vol. 13, no. 4, pp. 455–492, 1998.

[7] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," *arXiv preprint arXiv:0912.3995*, 2009.

[8] H. Liu, Y.-S. Ong, X. Shen, and J. Cai, "When gaussian process meets big data: A review of scalable gps," *IEEE transactions on neural networks and learning systems*, vol. 31, no. 11, pp. 4405–4423, 2020.

[9] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," in *Gaussian Processes in Practice*, Microsoft Research, April 2007.

[10] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *2011 IEEE International Conference on Robotics and Automation*, pp. 2845–2850, IEEE, 2011.

[11] J. Mahler, S. Patil, B. Kehoe, J. Van Den Berg, M. Ciocarlie, P. Abbeel, and K. Goldberg, "Gp-gpis-opt: Grasp planning with shape uncertainty using gaussian process implicit surfaces and sequential convex programming," in *2015 IEEE international conference on robotics and automation (ICRA)*, pp. 4919–4926, IEEE, 2015.

[12] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2015.

[13] H. J. Kushner, "A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise," 1964.

[14] J. Močkus, "On bayesian methods for seeking the extremum," in *IFIP Technical Conference on Optimization Techniques*, pp. 400–404, Springer, 1974.

[15] E. Brochu, V. M. Cora, and N. De Freitas, "A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," *arXiv preprint arXiv:1012.2599*, 2010.

[16] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.

[17] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2, pp. 235–256, 2002.

[18] R. M. Murray, Z. Li, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 2017.

[19] C. Ferrari, J. Canny, *et al.*, "Planning optimal grasps," in *Proceedings., 1992 IEEE International Conference on Robotics and Automation, 1992.*, vol. 3, pp. 2290–2295, IEEE, 1992.

[20] Z. Yi, R. Calandra, F. Veiga, H. van Hoof, T. Hermans, Y. Zhang, and J. Peters, "Active tactile object exploration with gaussian processes," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4925–4930, IEEE, 2016.

[21] D. Driess, P. Englert, and M. Toussaint, "Active learning with query paths for tactile object shape exploration," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 65–72, IEEE, 2017.

[22] G. Z. Gandler, C. H. Ek, M. Björkman, R. Stolkin, and Y. Bekiroglu, "Object shape estimation and modeling, based on sparse gaussian process implicit surfaces, combining visual data and tactile exploration," *Robotics and Autonomous Systems*, vol. 126, p. 103433, 2020.

[23] J. Nogueira, R. Martinez-Cantin, A. Bernardino, and L. Jamone, "Unscented bayesian optimization for safe robot grasping," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1967–1972, IEEE, 2016.

[24] J. Castanheira, P. Vicente, R. Martinez-Cantin, L. Jamone, and A. Bernardino, "Finding safe 3d robot grasps through efficient haptic exploration with unscented bayesian optimization and collision penalty," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1643–1648, IEEE, 2018.

[25] J. G. Lechuz-Sierra, A. E. H. Martin, A. M. Sundaram, R. Martinez-Cantin, and M. A. Roa, "Bayesian optimization for robust robotic grasping using a sensorized compliant hand," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 10503–10510, 2024.

[26] Y. Chen, A. E. Tekden, M. P. Deisenroth, and Y. Bekiroglu, "Sliding touch-based exploration for modeling unknown object shape with multi-fingered hands," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8943–8950, IEEE, 2023.

[27] C. De Farias, B. Tamadazte, M. Adjigble, R. Stolkin, and N. Marturi, "Task-informed grasping of partially observed objects," *IEEE Robotics and Automation Letters*, vol. 9, no. 10, pp. 8394–8401, 2024.

[28] M. K. Titsias, "Variational learning of inducing variables in sparse gaussian processes," in *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, vol. 5 of *Proceedings of Machine Learning Research*, pp. 567–574, 2009.

[29] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*, vol. 45, no. 1–3, pp. 503–528, 1989.

[30] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. A. Szerlip, P. Horsfall, and N. D. Goodman, "Pyro: Deep universal probabilistic programming," *CoRR*, vol. abs/1810.09538, 2018.

[31] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set," *IEEE Robotics amp; Automation Magazine*, vol. 22, no. 3, p. 36–52, 2015.

A. Detailed Experimental Results

This appendix provides the run-level results used to calculate the quantitative summaries reported in Chapter 5. For all quantitative analyses, the initial 50 warm-up evaluations, corresponding to iterations 0–49, are excluded. The original iteration numbering is retained, such that the first evaluation included in the analysis corresponds to iteration 50.

For each run, $t_{0.5}^{\mathrm{PFC}}$ and $t_{0.8}^{\mathrm{PFC}}$ denote the first iterations at which $P_{\mathrm{FC}} \geq 0.5$ and $P_{\mathrm{FC}} \geq 0.8$, respectively. The best PFC is the maximum PFC obtained during the analyzed part of the run, while the minimum surface cost is the lowest GPIS surface cost obtained over the same interval. A dash ("–") is used when a PFC threshold is not reached.

A.1. Proposed Geometry-Aware Acquisition

A.1.1. Last Restart Strategy

Table A.1 reports the individual results for the proposed geometry-aware acquisition using the `last` restart strategy. Results are shown for all five objects, all three exploration settings, and all five random seeds.

Table A.1.: Run-level results for the proposed geometry-aware acquisition using the `last` restart strategy.

Object	κ	Restart	Seed	$t_{0.5}^{\mathrm{PFC}}$	$t_{0.8}^{\mathrm{PFC}}$	Best PFC	Min. cost
Mustard bottle	$4.0 \rightarrow 0.8$	`last`	42	50	98	1.0	0.000

Continued on next page

Table A.1 – continued from previous page

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Mustard bottle	$4.0 \rightarrow 0.8$	last	2	55	130	0.9	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	5	52	79	1.0	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	7	55	72	1.0	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	17	59	62	1.0	0.000
Mustard bottle	2	last	42	52	66	1.0	0.000
Mustard bottle	2	last	2	58	58	1.0	0.000
Mustard bottle	2	last	5	61	88	1.0	0.000
Mustard bottle	2	last	7	55	68	1.0	0.000
Mustard bottle	2	last	17	54	124	1.0	0.000
Mustard bottle	3	last	42	67	84	1.0	0.000
Mustard bottle	3	last	2	59	75	1.0	0.000
Mustard bottle	3	last	5	79	79	1.0	0.000
Mustard bottle	3	last	7	61	82	1.0	0.000
Mustard bottle	3	last	17	52	52	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	42	62	127	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	2	61	152	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	5	54	125	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	7	52	75	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	17	82	168	1.0	0.000
Tuna fish can	2	last	42	51	146	1.0	0.000
Tuna fish can	2	last	2	65	152	1.0	0.000
Tuna fish can	2	last	5	92	143	1.0	0.000
Tuna fish can	2	last	7	74	74	1.0	0.000
Tuna fish can	2	last	17	51	109	1.0	0.000

Continued on next page

Table A.1 – continued from previous page

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Tuna fish can	3	last	42	68	100	1.0	0.000
Tuna fish can	3	last	2	61	61	1.0	0.000
Tuna fish can	3	last	5	65	76	0.9	0.000
Tuna fish can	3	last	7	53	101	1.0	0.000
Tuna fish can	3	last	17	57	139	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	last	42	103	143	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	last	2	62	74	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	last	5	56	56	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	last	7	50	50	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	last	17	51	51	1.0	0.000
Cracker box	2	last	42	95	142	1.0	0.000
Cracker box	2	last	2	58	58	1.0	0.000
Cracker box	2	last	5	60	60	1.0	0.000
Cracker box	2	last	7	78	164	1.0	0.000
Cracker box	2	last	17	52	52	1.0	0.000
Cracker box	3	last	42	98	98	1.0	0.000
Cracker box	3	last	2	111	144	1.0	0.000
Cracker box	3	last	5	74	75	1.0	0.000
Cracker box	3	last	7	54	54	1.0	0.000
Cracker box	3	last	17	52	52	1.0	0.000
Mini soccer ball	$4.0 \rightarrow 0.8$	last	42	53	53	1.0	0.000
Mini soccer ball	$4.0 \rightarrow 0.8$	last	2	51	65	1.0	0.000
Mini soccer ball	$4.0 \rightarrow 0.8$	last	5	50	66	1.0	0.000

Continued on next page

Table A.1 – continued from previous page

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Mini soccer ball	$4.0 \to 0.8$	last	7	60	62	1.0	0.000
Mini soccer ball	$4.0 \to 0.8$	last	17	62	75	1.0	0.000
Mini soccer ball	2	last	42	52	54	1.0	0.000
Mini soccer ball	2	last	2	51	51	1.0	0.000
Mini soccer ball	2	last	5	58	58	1.0	0.000
Mini soccer ball	2	last	7	50	50	1.0	0.000
Mini soccer ball	2	last	17	53	54	1.0	0.000
Mini soccer ball	3	last	42	51	56	1.0	0.000
Mini soccer ball	3	last	2	52	70	1.0	0.000
Mini soccer ball	3	last	5	54	54	1.0	0.000
Mini soccer ball	3	last	7	57	57	1.0	0.000
Mini soccer ball	3	last	17	55	57	1.0	0.000
Rubik's cube	$4.0 \to 0.8$	last	42	62	62	1.0	0.000
Rubik's cube	$4.0 \to 0.8$	last	2	53	53	1.0	0.000
Rubik's cube	$4.0 \to 0.8$	last	5	51	130	1.0	0.000
Rubik's cube	$4.0 \to 0.8$	last	7	64	95	1.0	0.000
Rubik's cube	$4.0 \to 0.8$	last	17	51	51	1.0	0.000
Rubik's cube	2	last	42	54	87	1.0	0.000
Rubik's cube	2	last	2	50	50	1.0	0.000
Rubik's cube	2	last	5	51	51	1.0	0.000
Rubik's cube	2	last	7	57	78	1.0	0.000
Rubik's cube	2	last	17	51	59	1.0	0.000
Rubik's cube	3	last	42	58	58	1.0	0.000

Continued on next page

Table A.1 – continued from previous page

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Rubik's cube	3	last	2	53	54	1.0	0.000
Rubik's cube	3	last	5	51	51	1.0	0.000
Rubik's cube	3	last	7	52	102	1.0	0.000
Rubik's cube	3	last	17	50	95	1.0	0.000

A.1.2. Best Restart Strategy

Table A.2 reports the run-level results for the `best` restart strategy used in the restart-strategy comparison. Results are shown for the mustard bottle, tuna fish can, and cracker box using the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting and fixed $\kappa = 2$.

Table A.2.: Run-level results for the proposed geometry-aware acquisition using the `best` restart strategy.

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Mustard bottle	$4.0 \rightarrow 0.8$	best	42	50	79	1.0	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	best	2	55	207	1.0	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	best	5	52	138	0.9	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	best	7	55	72	1.0	0.000
Mustard bottle	$4.0 \rightarrow 0.8$	best	17	59	62	0.9	0.000
Mustard bottle	2	best	42	52	66	1.0	0.000
Mustard bottle	2	best	2	58	58	1.0	0.000
Mustard bottle	2	best	5	61	134	1.0	0.000
Mustard bottle	2	best	7	72	72	1.0	0.000
Mustard bottle	2	best	17	54	114	1.0	0.000

Continued on next page

Table A.2 – continued from previous page

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Tuna fish can	$4.0 \rightarrow 0.8$	best	42	62	128	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	best	2	61	124	0.9	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	best	5	78	104	1.0	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	best	7	52	111	0.9	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	best	17	82	109	1.0	0.000
Tuna fish can	2	best	42	51	97	1.0	0.000
Tuna fish can	2	best	2	61	107	1.0	0.000
Tuna fish can	2	best	5	78	169	1.0	0.000
Tuna fish can	2	best	7	74	139	1.0	0.000
Tuna fish can	2	best	17	51	107	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	best	42	102	102	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	best	2	62	74	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	best	5	56	56	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	best	7	50	50	1.0	0.000
Cracker box	$4.0 \rightarrow 0.8$	best	17	51	51	1.0	0.000
Cracker box	2	best	42	95	141	1.0	0.000
Cracker box	2	best	2	58	58	1.0	0.000
Cracker box	2	best	5	60	60	1.0	0.000
Cracker box	2	best	7	78	133	1.0	0.000
Cracker box	2	best	17	52	52	1.0	0.000

A.2. Method Based on de Farias et al.

Table A.3 reports the individual run-level results for the method based on the combined GPIS–PFC formulation of de Farias et al. Results are shown for the mustard bottle, tuna fish can, and cracker box using the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting and fixed $\kappa = 2$. Each configuration was evaluated using five random seeds.

Table A.3.: Run-level results for the method based on the combined GPIS–PFC formulation of de Farias et al.

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Mustard bottle	$4.0 \rightarrow 0.8$	last	42	–	–	0.0	1.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	2	–	–	0.0	1.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	5	–	–	0.0	1.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	7	–	–	0.0	1.000
Mustard bottle	$4.0 \rightarrow 0.8$	last	17	–	–	0.0	1.000
Mustard bottle	2	last	42	–	–	0.0	1.000
Mustard bottle	2	last	2	–	–	0.0	0.750
Mustard bottle	2	last	5	–	–	0.0	1.000
Mustard bottle	2	last	7	–	–	0.0	0.750
Mustard bottle	2	last	17	–	–	0.0	0.500
Tuna fish can	$4.0 \rightarrow 0.8$	last	42	–	–	0.0	1.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	2	–	–	0.0	1.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	5	–	–	0.0	1.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	7	–	–	0.0	1.000
Tuna fish can	$4.0 \rightarrow 0.8$	last	17	–	–	0.0	1.000
Tuna fish can	2	last	42	–	–	0.0	1.000
Tuna fish can	2	last	2	–	–	0.0	0.267

Continued on next page

Table A.3 – continued from previous page

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Tuna fish can	2	last	5	–	–	0.0	0.500
Tuna fish can	2	last	7	–	–	0.0	0.267
Tuna fish can	2	last	17	–	–	0.0	0.500
Cracker box	$4.0 \rightarrow 0.8$	last	42	–	–	0.0	1.000
Cracker box	$4.0 \rightarrow 0.8$	last	2	–	–	0.0	1.000
Cracker box	$4.0 \rightarrow 0.8$	last	5	–	–	0.0	1.000
Cracker box	$4.0 \rightarrow 0.8$	last	7	–	–	0.0	1.000
Cracker box	$4.0 \rightarrow 0.8$	last	17	–	–	0.0	1.000
Cracker box	2	last	42	–	–	0.0	0.667
Cracker box	2	last	2	–	–	0.0	0.509
Cracker box	2	last	5	–	–	0.0	0.478
Cracker box	2	last	7	–	–	0.0	0.672
Cracker box	2	last	17	–	–	0.0	0.340

A.3. Hand-Inspired Contact Configuration

Table A.4 reports the individual run-level results for the hand-inspired version of the proposed geometry-aware acquisition. The experiment uses the adaptive $\kappa : 4.0 \rightarrow 0.8$ setting and the `last` restart strategy on the mustard bottle, tuna fish can, and cracker box.

Table A.4.: Run-level results for the hand-inspired contact configuration experiment.

Object	κ	Restart	Seed	$t_{0.5}^{PFC}$	$t_{0.8}^{PFC}$	Best PFC	Min. cost
Mustard bottle	$4.0 \to 0.8$	`last`	42	124	141	1.0	0.000
Mustard bottle	$4.0 \to 0.8$	`last`	2	56	56	1.0	0.000
Mustard bottle	$4.0 \to 0.8$	`last`	5	100	100	1.0	0.000
Mustard bottle	$4.0 \to 0.8$	`last`	7	56	106	1.0	0.000
Mustard bottle	$4.0 \to 0.8$	`last`	17	74	109	1.0	0.000
Tuna fish can	$4.0 \to 0.8$	`last`	42	74	–	0.6	0.000
Tuna fish can	$4.0 \to 0.8$	`last`	2	64	115	1.0	0.000
Tuna fish can	$4.0 \to 0.8$	`last`	5	71	76	1.0	0.000
Tuna fish can	$4.0 \to 0.8$	`last`	7	75	320	0.9	0.000
Tuna fish can	$4.0 \to 0.8$	`last`	17	85	141	0.9	0.000
Cracker box	$4.0 \to 0.8$	`last`	42	59	82	1.0	0.000
Cracker box	$4.0 \to 0.8$	`last`	2	52	52	1.0	0.000
Cracker box	$4.0 \to 0.8$	`last`	5	72	72	1.0	0.000
Cracker box	$4.0 \to 0.8$	`last`	7	60	60	1.0	0.000
Cracker box	$4.0 \to 0.8$	`last`	17	68	74	1.0	0.000

A.4. Extended Surface-Based Warm-Up

Table A.5 reports the individual run-level results for the extended surface-based warm-up experiment. The experiment was performed on the mustard bottle and tuna fish can using the adaptive $\kappa : 4.0 \to 0.8$ setting and fixed $\kappa = 2$, with the `last` restart strategy and five random seeds per configuration.

For this experiment, the first 100 evaluations correspond to the surface-based warm-up. Accordingly, the PFC thresholds, best PFC values, and minimum GPIS surface costs

reported in this section are evaluated from iteration 100 onward. The original iteration numbering is retained.

Table A.5.: Run-level results for the extended surface-based warm-up experiment.

Object	κ	Restart	Seed	$t_{0.5}^{\mathrm{PFC}}$	$t_{0.8}^{\mathrm{PFC}}$	Best PFC	Min. cost
Mustard bottle	$4.0 \rightarrow 0.8$	`last`	42	–	–	0.0	0.251
Mustard bottle	$4.0 \rightarrow 0.8$	`last`	2	–	–	0.0	0.500
Mustard bottle	$4.0 \rightarrow 0.8$	`last`	5	–	–	0.0	0.250
Mustard bottle	$4.0 \rightarrow 0.8$	`last`	7	–	–	0.0	0.250
Mustard bottle	$4.0 \rightarrow 0.8$	`last`	17	–	–	0.0	0.250
Mustard bottle	2	`last`	42	–	–	0.0	0.192
Mustard bottle	2	`last`	2	–	–	0.0	0.217
Mustard bottle	2	`last`	5	–	–	0.0	0.253
Mustard bottle	2	`last`	7	–	–	0.0	0.252
Mustard bottle	2	`last`	17	–	–	0.0	0.015
Tuna fish can	$4.0 \rightarrow 0.8$	`last`	42	258	259	0.9	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	`last`	2	275	337	0.8	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	`last`	5	273	324	0.8	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	`last`	7	261	336	0.8	0.000
Tuna fish can	$4.0 \rightarrow 0.8$	`last`	17	281	–	0.5	0.000
Tuna fish can	2	`last`	42	162	203	0.9	0.000
Tuna fish can	2	`last`	2	144	–	0.5	0.000
Tuna fish can	2	`last`	5	139	314	0.8	0.000
Tuna fish can	2	`last`	7	141	–	0.6	0.000
Tuna fish can	2	`last`	17	145	–	0.7	0.000