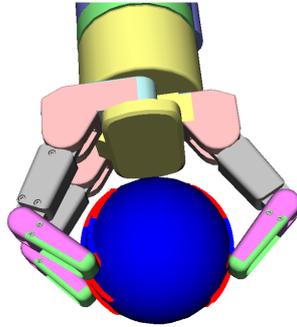




INSTITUTO SUPERIOR TÉCNICO
Universidade Técnica de Lisboa



Robotic Grasp Optimization from Contact Force Analysis

Filipe Fernandes Veiga

Dissertation submitted for obtaining the degree of
Master in Electrical and Computer Engineering

Jury

President: Carlos Filipe Gomes Bispo
Supervisor: Alexandre José Malheiro Bernardino
Co-Supervisor: José Alberto Rosado dos Santos Victor
Member: João Fernando Cardoso Silva Sequeira

April 2012

Science may set limits to knowledge, but should not set limits to imagination.

Bertrand Russell

Acknowledgments

I would like to thank my supervisor Alexandre Bernardino for the help and tireless support during the course of this work. Secondly I would like to thank my family, specially my mother Isabel Fernandes, my father Manuel Fernando Veiga, my grandmother Irene Macedo Moura and my grandfather Mario Jesus Fernandes for all the advice, support and motivation that made me surpass all the challenges brought on by life to date. I would like to thank Inês Marques Dias for the love she shows me everyday.

Abstract

This thesis addresses the problem of robotic grasp optimization. Due to uncertainty both on the robot kinematics, control and perception, it is very hard to analytically compute good grasps and execute them successfully. Our approach is based on searching for the best grasp configurations by iteratively optimizing a suitable grasp criterion. This approach may be compared to human learning stages where infants learn by trial and error what are the best grasping strategies. Initial grasps are often unsuccessful, but after a few trials the system learns to adapt to the uncertainties in the environment.

The thesis is composed by two main parts. On the first part we study the problem of grasp quality assessment. This is an essential ability both in off-line planning of grasps for robotic manipulation as in evaluating on-line executed grasps. Until now existing metrics rely on force-Closure tests that are limited in scope. We propose a grasp quality metric that provides coherent performance criteria even for non-force-Closure grasps. It is based in two improvements with respect to existing metrics: a surface contact model, providing smoothness properties, and a bimodal wrench space analysis that provides continuity in the transition from non-force-closure to force-closure grasps. It is experimentally demonstrated, in a series of simulations, that the developed metric outperforms existing ones. We also show empirically that optimal points are invariant to the choice of reference frame, which allows its application on-line in non-model based approaches. On the second part of the thesis we study the problem of efficiently searching for the best grasp. Naive optimization of the grasp quality metric (e.g. uniform sampling) becomes intractable very quickly as the search space parameters grow. To address this problem we apply recent methods on Bayesian Optimization to obtain good approximations to global optima in a reduced number of trials. Gaussian Processes are used to encode our knowledge and uncertainty of the grasp quality metric as new trials are performed. Then, new samples are planned by searching in the Gaussian Process for points with favorable expected improvement. This last search is performed using recent Global Optimization methods.

We present results both in simulation and on a real robotic platform in the context of the HANDLE project. Results show clear benefits of our approach both in what respects the ability of the grasp quality assessment metric to represent enlarged parts of the search space and in what respects the

ability of Bayesian Optimization to reduce the number of trials needed to find the optimal grasps.

Keywords

Grasping, Optimization, Robot Hand, Evaluation Metric, Contact Model, Wrench Analysis.

Resumo

Esta tese aborda o problema da optimização de um agarre robótico. Devido à incerteza tanto na cinemática do robot como nos seus sistemas de controlo e percepção, o calculo analítico de bons agarres robóticos torna-se num problema difícil. A nossa abordagem ao problema baseia-se na procura iterativa das melhores configurações de agarre, optimizando um determinado critério. Esta abordagem pode ser comparada com algumas etapas de aprendizagem humana onde crianças aprendem as melhores estratégias de agarre por tentativa e erro. Apesar dos agarres iniciais não serem bem sucedidos, com o decorrer dos ensaios o sistema aprende a adaptar-se às incertezas.

A tese é composta por duas partes principais. Na primeira parte estuda-se o problema de como avaliar a qualidade de um agarre. Esta é uma capacidade essencial tanto em planeamento “off-line” de agarres para manipulação robótica como na avaliação “on-line” de agarres executados. Até agora métricas existentes dependiam de testes de fecho-de-força, o que limitava o espaço de agarres que conseguiam avaliar. Propõe-se uma métrica de qualidade para agarres robóticos que fornece critérios de avaliação sensíveis às propriedades de todos os agarres, mesmo quando estes não passam o teste de fecho-de-força. Esta métrica baseia-se em duas melhorias em relação às métricas já existentes: um modelo de superfície de contacto que proporciona uma representação mais realista do que se passa ao nível do contacto entre objecto e manipulador robótico e uma análise bimodal do espaço força/torque que assegura transições continuas entre agarres fecho-de-força e agarres não fecho-de-força. É demonstrado experimentalmente, através de simulação, que a métrica desenvolvida supera as existentes. Mostra-se empiricamente que configurações óptimas são invariantes à escolha do referencial de referência, o que permite a utilização do sistema em aplicações “on-line” e aplicações sem modelo para o objecto alvo. Na segunda parte da tese, focamos-nos na procura eficiente do agarre ideal. Uma optimização ingénua da métrica de qualidade (e.g. amostragem uniforme) torna-se rapidamente impraticável com o aumento do espaço de parâmetros. Para abordar este problema recorreremos a métodos de optimização Bayesianos por forma a obter uma boa aproximação do valor óptimo global num numero reduzido de ensaios. Processos Gaussianos são utilizados para codificar o nosso conhecimento e a incerteza acerca da métrica de qualidade à medida que se fazem novos ensaios. Depois, novos ensaios são planeados procurando pontos no Processo Gaussiano que forneçam valores favoráveis para a melhoria esperada. Esta última procura é realizada com o auxilio de métodos recentes de optimização global.

São apresentados resultados tanto em ambiente de simulação como numa plataforma real no contexto do projecto HANDLE. Os resultados demonstram claros benefícios na utilização da nossa

abordagem tanto no que toca às propriedades da métrica de qualidade de agarres como no que diz respeito à capacidade da otimização Bayesiana de reduzir o número de ensaios necessários para encontrar agarres ótimos.

Palavras Chave

Agarre Robótico, Otimização, Mão Robótica, Métrica de Avaliação, Modelo de Contacto, espaço força/torque

Contents

1	Introduction	1
1.1	Motivation	2
1.2	State of The Art	4
1.3	Original Contributions	4
1.4	Thesis Outline	5
2	Grasp Quality Metric	7
2.1	Wrenches	8
2.2	Contact Model	9
2.3	Contact Surface Model	11
2.4	Grasp Representation	12
2.5	Closure	13
2.6	Bimodal Wrench Space Analysis	14
3	Bayesian Optimization	17
3.1	Gaussian Process Regression	18
3.2	Expected Improvement	19
3.3	Direct Optimization Algorithm	21
3.4	The Complete Method	22
4	Experimental Results	25
4.1	Experimental Setup	26
4.2	Metric Comparisons	27
4.3	Contact Model Influence	29
4.4	Reference Frame Independence	30
4.5	Searching with two parameters	31
4.6	1D exploration with Bayesian Optimization	34
4.7	2D exploration with Bayesian Optimization	36
4.8	Bayesian optimization versus random sampling	40
4.9	Results on a real platform	42
5	Conclusions and Future Work	45

List of Figures

1.1	The Shadow Robot Hand equipped with fingertip force/torque sensors ATI NANO 17.	3
2.1	Frictionless point contact model.	9
2.2	Point contact with friction model.	10
2.3	Soft-finger contact model.	11
2.4	Steps for generating the Surface Contact Model. Line segments generated (top left), Line segments and their intersections with the object surface (top right) and intersection points and their correspondig surface normals (bottom). The blue box represents the object, the red points represent the intersections between the line segments and the object and the light green lines represent the object surface normal at each intersection point.	12
2.5	Convex Hull of the grasp wrenchs, the Grasp Wrench Space	13
2.6	The Largest Sphere metric. Used by the BW metric to analyse force-closure grasps.	14
2.7	The BW metric. The grasp wrench space distance to force-closure when analysing a non-force-closure grasp.	15
3.1	Flowchart depicting the complete Bayesian Optimization method.	23
4.1	2D sampling of a Sphere	26
4.2	The experimental setup used for validating the proposed metric.	27
4.3	Comparison between Largest Sphere Metric (top), Grasp Wrench Volume Metric (mid- dle) and BW Metric (bottom) for grasps on a sphere.	27
4.4	Comparison between Largest Sphere Metric (top), Grasp Wrench Volume Metric (mid- dle) and BW Metric (bottom) for grasps on a rotated cuboid.	28
4.5	Comparison between Largest Sphere Metric (top), Grasp Wrench Volume Metric (mid- dle) and BW Metric (bottom) for grasps on a star prism.	29
4.6	BW Metric values using diferent contact models when executing grasps on a cuboid (left) and on a sphere (right).	29
4.7	BW Metric values of grasps in a sphere, computed in a grid sampling of the hand position parameters with the mean point model (left) and with the projected patch model (right).	30

4.8	Metric Results using as reference frame the World frame and the Object frame (left) and some of the robot's frames (right)	30
4.9	<i>BW</i> Metric values of grasps in a sphere, computed in a grid sampling of the hand position parameters.	31
4.10	<i>BW</i> Metric values of grasps in a wine glass, computed in a grid sampling of the hand position parameters.	31
4.11	<i>BW</i> Metric values of grasps in a cylinder, computed in a grid sampling of the hand position parameters.	32
4.12	<i>BW</i> Metric values of grasps in a mug, computed in a grid sampling of the hand position parameters.	32
4.13	<i>BW</i> Metric values of grasps in a cuboid, computed in a grid sampling of the hand position parameters.	33
4.14	<i>BW</i> Metric values of grasps in a rotated cuboid, computed in a grid sampling of the hand position parameters.	33
4.15	<i>BW</i> Metric values of grasps in a star prism, computed in a grid sampling of the hand position parameters.	34
4.16	1D scan of a sphere, computed by sampling along the input parameter.	35
4.17	1D scan of a star prism, computed by sampling along the input parameter.	36
4.18	<i>BW</i> Metric values of grasps in a sphere, computed using the Bayesian Optimization method.	37
4.19	<i>BW</i> Metric values of grasps in a wine glass, computed using the Bayesian Optimization method.	37
4.20	<i>BW</i> Metric values of grasps in a cylinder, computed using the Bayesian Optimization method.	38
4.21	<i>BW</i> Metric values of grasps in a mug, computed using the Bayesian Optimization method.	38
4.22	<i>BW</i> Metric values of grasps in a cuboid, computed using the Bayesian Optimization method.	39
4.23	<i>BW</i> Metric values of grasps in a rotated cuboid, computed using the Bayesian Optimization method.	39
4.24	<i>BW</i> Metric values of grasps in a star prism, computed using the Bayesian Optimization method.	40
4.25	Comparing the evolution of best value found when sampling a sphere with the Bayesian approach versus the random sampling method.	41
4.26	Comparing the evolution of best value found when sampling a mug with the Bayesian approach versus the random sampling method.	41
4.27	Comparing the evolution of best value found when sampling a star prism with the Bayesian approach versus the random sampling method.	42
4.28	Shadow arm and hand that serves as main platform for the HANDLE project (left), image captured by the Kinect mounted on the arm during the course of the experiment.	42

4.29 Metric values taken during the experiment. 43

List of Tables

4.1	Number of Samples taken by each method for each object.	40
-----	---	----

List of Symbols

\mathbb{R}	Set of Real Numbers.
\mathbf{f}	Force.
$\boldsymbol{\tau}$	Torque.
$\boldsymbol{\omega}$	Wrench.
R_{cb}	Rotation matrix from coordinate frame b to coordinate frame c.
$[\mathbf{p}]_x$	Skew-Symmetric matrix of \mathbf{p} .
μ_f	Static Friction Coefficient.
γ	Torsional Friction Coefficient.
f^t	Magnitude of the tangential force component.
f^n	Magnitude of the normal force component.
ω_{c_i}	Set of wrenches generated by contact i .
\mathbf{f}_{c_i}	Force magnitudes exerted by contact i .
\mathbf{p}_m	Mean position of the contact point cloud for a single finger.
R_m	Rotation matrix from frame centered at c_m to the global contact reference frame.
G_i	Set of wrenches generated by contact i expressed on the global contact reference frame.
G	Grasp Map.
W_G	Grasp Wrench Space.
W_0	Wrench space origin.
\mathcal{GP}	Gaussian Process.
μ	\mathcal{GP} mean function.
Σ	\mathcal{GP} covariance matrix.
K	Kernel.
\mathbf{X}	Observation point set.
\mathbf{X}_*	Estimation point set.
\mathbf{F}_*	Estimated function values point set.
\mathbf{F}	Observed function values point set.
σ_n^2	Observation noise variance.
I	Improvement function.
\mathcal{D}_n	Knowledge at time n.
EI	Expected Improvement function.
f^{max}	Current function maximum value.
$\hat{\sigma}_f^2$	Estimated function variance.
$\bar{\Omega}$	Direct hypercube domain.
\hat{K}	Lipchitz Constant.
f_{min}	Current function minimum value.
d_j	Dimensional measure of hypercube j .

1

Introduction

Contents

1.1 Motivation	2
1.2 State of The Art	4
1.3 Original Contributions	4
1.4 Thesis Outline	5

1.1 Motivation

Robotic grasping is a fundamental skill in industrial and service robots. Most modern robots have manipulators either in the form of simple grippers or highly dexterous robotic hands. However, the execution of reliable grasps in objects has been a topic of puzzling difficulty. In industrial environments it is already common to see robots performing repetitive tasks with very simple manipulators. In these cases the robots perform on controlled environments and are equipped with specific task oriented manipulators depending on their purpose.

The amount of variables that are involved in the grasping problem for more complex manipulators make grasp planning hardly tractable and execution of grasps brittle and very prone to failures. Uncertainty in the measurement of position and shape of the object, unmodeled friction characteristics and imperfect object models, are limiting the amount of manipulation tasks that current robots can perform autonomously. The main difficulty relies on measuring with precision the state of the object and hand in terms of the physical variables that influence a grasp experience - position and velocities of the object, contact points, forces, friction, slip, softness, etc.

Common setups for robotic manipulation are composed by a robotic arm-hand system and a set of sensors to detect the shape and position of objects of interest (typically vision and/or 3D range sensors), usually calibrated with respect to the manipulator reference system. It is also common nowadays to have hands instrumented with some form of tact sensing to determine contact with objects. The typical approach to grasping is composed by three main stages: first the external sensors are used to obtain an estimate of the object pose and a desired position for the manipulator is computed by a planner module; then a control algorithm drives the hand to the desired position; and finally the hand fingers are closed until pressure is exerted in the object. However, after the fingers are closed, the process stops and no further grasp improvement is followed. Due to multiple uncertainty sources, fingers may be misplaced in the object and it may slip or fall down after being lifted. Because all decision is concentrated in the first phase of the manipulation, errors that accumulate during the control phase are not compensated and the grasp may fail. Better contact and force measurement, together with better grasp quality assessment criteria, are essential to compensate for the uncertainty in real-world grasp execution.

Also to consider are off-line grasp planning applications. Popular grasp simulators [1], [2] include grasp planners that generate thousands of hand configurations around the object to grasp. After one of these configurations is chosen, the fingers are closed until contact with the object is achieved. Contact points and forces are computed and a grasp quality metric evaluates the feasible grasp candidates. These candidates are obtained through uniform sampling of the hand configuration around the object. In these settings, reducing the number of required samples, through the use of smart search strategies, e.g. Bayesian optimization [3], leads to faster convergence to optimal grasp configurations.

This thesis addresses the computation of globally optimal grasping configurations using an incremental optimization approach. Initial grasps are planned in a very coarse manner based on object rough position and shape. These grasps are then evaluated and used to bootstrap new trials that will

improve the grasp quality. The approach can be used in two modalities: (i) in a simulator to plan optimal grasps, if the object is known and there is low uncertainty in the system or (ii) on-line in the real robot, when the object is unknown and there are perception uncertainties and calibration errors. Two components are essential for the approach: (i) a suitable grasp quality metric to assess the stability of large ranges of grasp configurations and able to operate with known and unknown objects; and (ii) an efficient global optimization method to search for the solution in a reduced number of trials.

The first part addresses the computation of grasp stability metrics to evaluate a grasp state and compare among different grasps. In particular we propose a metric that provides quantitative information also for non-stable grasps, that other metrics often disregard, thus largely extending the domain of application. This is a very important characteristic if one wants to search for better grasps locally via small changes in the hand parameters. After evaluating the properties of the metric with respect to the classical ones, the focus of the work shifts, emphasizing grasp optimization through active learning. Bayesian Optimization methods will be used to learn how to grasp an object in an efficient manner. It will be shown that our approach is valid for model or non-model based systems (where the object model is not known) and that for robotic hands with the sensors required to measure contact points and forces, it will be possible to optimize grasps by searching on hand parameter space. Even if not all robotic hands have such sensing abilities, it is becoming more common to instrument the fingertips with advanced tactile devices. For instance, the Shadow Robot Hand¹ has been recently equipped with force/torque sensors in each fingertip in a special encapsulation (see Fig. 1.1) that allows the reconstruction of the contact points and force information [4].



Figure 1.1: The Shadow Robot Hand equipped with fingertip force/torque sensors ATI NANO 17.

¹<http://www.shadowrobot.com/>

1.2 State of The Art

A notable work addressing grasp improvement is [5], where a set of rules derived from contact sensory feedback are applied during the reaching and finger closing phases of a manipulation trial. It uses the wrist force/torque sensor and the fingertip integrated force sensors of the Barrett hand, to implement a series of discrete correction stages to make sure a proper grasp stability is achieved.

Force-closure criteria [6] is one of the most used grasp evaluation criteria. In its basic form it makes a binary decision about the stability of a grasp: if the fingers can exert forces and torques in any direction then the grasp is force-closure. It means that, if infinite forces are allowed, the grasp could resist to any external wrenches. This criteria has been refined in several ways to provide continuous measures of grasp stability. The Grasp Wrench Volume metric has been proposed in [7] to evaluate grasps. The larger the volume of the grasp wrench, the more stable the grasp will be, in qualitative terms. The Largest Sphere criterion [8] computes the magnitude of the largest worst-case disturbance that can be resisted by a grasp of unit strength in any direction. Seeing that the grasping field is quite a large field, refer to [9] for a more complete review on grasping.

Only recently have the scientific community turned to Bayesian Optimization methods as a means of solving Global Optimization problems. Recently applied to machine learning problems such as location, path planning and hierarchical system control, Bayesian optimization methods are seen as one of the most efficient tools for minimizing the number of function evaluations required, [3].

1.3 Original Contributions

In this section we emphasize the contributions introduced by this work.

Surface contact model A model for the contact surface between a rigid body and a robotic gripper that allows a more realistic representation of the the contact points improving the overall performance of any metric with analysis based on these points.

Bimodal Wrench Space Analysis Metric A new grasp quality assessment metric that allows previously unevaluable non-force-Closure grasps to be evaluated. The evaluation of non-force-Closure grasp gives a sense of distance to force-Closure allowing optimization and control systems to react accordingly.

Application of Bayesian Optimization to the Grasping Problem This work presents an application of Bayesian Optimization to the grasping problem with on-line quality assessment, which is not common in the literature.

1.4 Thesis Outline

This thesis is composed of 5 chapters and is organized as follows. In Chapter 2 we present the formulation and the algorithm to compute the proposed grasp quality evaluation metric. In Chapter 3 the methods that allow the robot to learn efficiently are described. Chapter 4 presents the simulation setup and shows results illustrating the performance of the metric as well as the results of the learning procedure. Finally, Chapter 5 concludes our work, discussing possible applications and pointing directions for further research.

2

Grasp Quality Metric

Contents

2.1 Wrenches	8
2.2 Contact Model	9
2.3 Contact Surface Model	11
2.4 Grasp Representation	12
2.5 Closure	13
2.6 Bimodal Wrench Space Analysis	14

The modern approach to robotic grasping tries to understand and emulate how humans use their hands to explore, restrain and manipulate objects. On account of the human hand complexity which provides extensive sensory feedback (sensing slip, object weight, object stability, etc.), humans intuitively evaluate the quality of the grasp they are performing. Robots on the other hand have manipulator and sensor limitations that increase the difficulty of the grasping process. A quality metric for grasp assessment is required as an interpretation of the sensory feedback given by the robot manipulator. Measuring the quality of stable grasps but also giving a measure of the distance to stability of non-stable grasps, while transitioning from stable to non-stable grasps in a smooth fashion, are the core aspects sought-after when designing a grasp quality metric.

Assuring that the metric is designed to fulfill the above requirements is not enough do ensure precise grasp assessment. As the metric's inputs, the contact points are crucial to the performance of the metric and achieving a realistic representation of these contacts leads to significant error reduction and more accurate results.

This chapter will depict how it is possible to attain a fair reproduction of what is contemplated on the real system by using a realistic model of the contact points in conjunction with a bimodal wrench space analysis metric. After starting with a basic description of all the involved basic concepts, the concept of contact surface is introduced. With the models of the contacts it is possible to define a grasp and outline some important grasp closure properties. Finally the chapter comes to a close with a step by step walk-through of the wrench space analysis.

2.1 Wrenches

A grasp is no more than a set of forces applied on a rigid body by a manipulator. Each of these forces consists of a linear component (pure force) and an angular component (pure moment) acting at a point. Representing this force/moment pair as a vector in \mathbb{R}^6 defines a wrench w , [6].

$$w = \begin{bmatrix} \mathbf{f} \\ \boldsymbol{\tau} \end{bmatrix} \quad \begin{array}{l} \mathbf{f} \in \mathbb{R}^3 \\ \boldsymbol{\tau} \in \mathbb{R}^3 \end{array} \quad (2.1)$$

The values of this wrench vector $w \in \mathbb{R}^6$ depend on the coordinate frame in which the force and the moment are represented. If B is a coordinate frame attached to a rigid body, then $w_b = (\mathbf{f}_b, \boldsymbol{\tau}_b)$ is the wrench applied at the origin of B, with \mathbf{f}_b and $\boldsymbol{\tau}_b$ specified with respect to the B coordinate frame.

If several wrenches are applied on the same rigid body the resulting net wrench can be constructed by adding all the wrench vectors. This addition only makes sense if all the wrench vectors are represented with respect to the same coordinate frame. Thus, if all wrenches in a wrench set w_i are to be added, all the wrench vectors must be rewritten to a single coordinate frame before the addition is performed. If B and C are distinct coordinate frames, the transformation of a wrench w_b applied at the origin of the B frame to an equivalent wrench w_c applied at the origin of the C frame can be achieved by

$$\begin{bmatrix} \mathbf{f}_c \\ \boldsymbol{\tau}_c \end{bmatrix} = \begin{bmatrix} R_{cb} & 0 \\ [\mathbf{p}_{bc}]_{\times} R_{cb} & R_{cb} \end{bmatrix} \begin{bmatrix} \mathbf{f}_b \\ \boldsymbol{\tau}_b \end{bmatrix}. \quad (2.2)$$

where $[\mathbf{p}_{bc}]_{\times}$ is the skew-symmetric matrix of \mathbf{p}_{bc} defined by

$$[\mathbf{p}]_{\times} = \begin{bmatrix} 0 & -p_3 & p_2 \\ p_3 & 0 & -p_1 \\ -p_2 & p_1 & 0 \end{bmatrix} \quad (2.3)$$

where p_1 , p_2 and p_3 are the components of vector \mathbf{p} .

With the use of this matrix the transformation includes an additional torque $\mathbf{p}_{bc} \times \mathbf{f}_c$ which is the torque generated by applying the force \mathbf{f}_b at \mathbf{p}_{bc} .

2.2 Contact Model

Using wrenches as a representation of forces and torques, it now becomes necessary to model the contact between the object and the manipulator. This model will provide the wrenches produced at each contact point. Each contact point is represented by a coordinate frame, R_{c_i} , attached to the contact location, \mathbf{p}_{c_i} , which is defined by its relative position and orientation with respect to a reference frame R_r . For example, the coordinate frame centered on the object's center of mass. The coordinate frame R_{c_i} is chosen such that its z-axis points in the direction of the surface normal at the point of contact.

Several models for these contact points have been considered. The frictionless point contact model, [6], is the simplest of the considered models. In this model no friction between the contacting surfaces is taken into account. Thus the model only allows forces along the surface normal direction. Using this model the produced wrenches at each contact c_i can be represented as

$$w_{c_i} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} f_{c_i}, \quad f_{c_i} \geq 0 \quad (2.4)$$

where $f_{c_i} \in \mathbb{R}$ is the magnitude of the applied normal force.

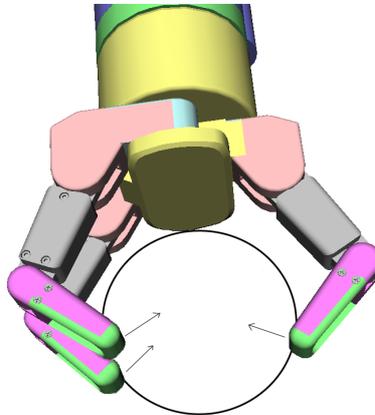


Figure 2.1: Frictionless point contact model.

Since not considering friction is a fairly naive approach, a friction model has to be introduced. The friction model chosen was Coulomb's friction model, [6]. This model states that forces in the tangential direction to the contact surface can be introduced with maximum magnitude proportional to the magnitude of the normal component. The constant of proportionality between the two components is the static coefficient of friction μ_f

$$|f^t| \leq \mu_f |f^n| \quad (2.5)$$

where $f^t \in \mathbb{R}$ and $f^n \in \mathbb{R}$ are the magnitudes of the tangential and normal force components respectively. The geometric meaning of this condition is that any applied force has to lie inside a cone centered about the surface normal at the point of contact. This cone is called the friction cone. Using Coulomb's model, it is a simple task to derive the point contact with friction model.

$$w_{c_i} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \mathbf{f}_{c_i}, \quad \mathbf{f}_{c_i} \in FC_{c_i} \quad (2.6)$$

where the friction cone FC_{c_i} is

$$FC_{c_i} = \{\mathbf{f} \in \mathbb{R}^3 : \sqrt{f_1^2 + f_2^2} \leq \mu_f f_3, f_3 \geq 0\}. \quad (2.7)$$

The components f_1 , f_2 and f_3 represent the the force along the x , y and z coordinates respectively.

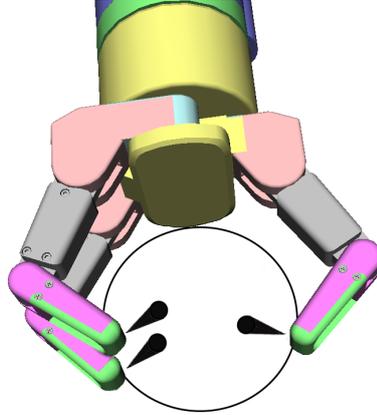


Figure 2.2: Point contact with friction model.

The soft-finger contact model, [6], is a more complete model that allows forces to be applied in a cone about the surface normal but also allows torques about that normal. Using this model the resulting wrench at each contact point is

$$w_{c_i} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{f}_{c_i}, \quad \mathbf{f}_{c_i} \in FC_{c_i} \quad (2.8)$$

where the friction cone FC_{c_i} is defined as

$$FC_{c_i} = \{\mathbf{f} \in \mathbb{R}^4 : \sqrt{f_1^2 + f_2^2} \leq \mu_f f_3, f_3 \geq 0, |f_4| \geq \gamma |f_3|\} \quad (2.9)$$

with γ representing the torsional friction coefficient and f_4 the torque magnitude along the contact normal direction. This was the chosen model for this work since it is the most realist model of the ones presented previously.

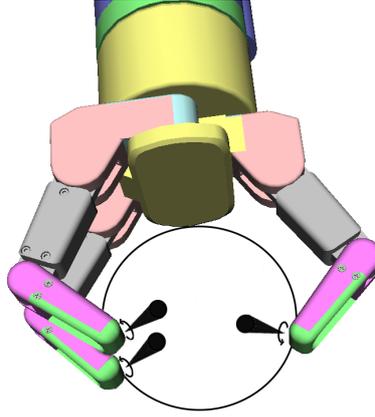


Figure 2.3: Soft-finger contact model.

While analyzing a grasp, the normal force is assumed to be of unit magnitude

$$f_3 = 1 \quad (2.10)$$

and the friction cone is sampled over its outer limits. The resulting set of wrenches at each contact point

$$w_{c_i}^k = \begin{bmatrix} \mu_f \sin(\theta_k) \\ \mu_f \cos(\theta_k) \\ 1 \\ 0 \\ 0 \\ \pm\gamma \end{bmatrix}, \quad \theta_k \in [0, 2\pi] \quad (2.11)$$

offer the base framework for grasp formulation.

One important aspect of this framework is that although the common choice for the reference frame R_r is the coordinate frame centered on the object's center of mass, the end results will be independent of the chosen reference frame, as will be shown in the experimental results. This allows the use of this framework on non-model based systems as long as, when comparing two grasps, both of them are represented with respect to the same reference.

2.3 Contact Surface Model

Using a single-point contact model might not be a completely realistic approach for contact modeling when using anthropomorphic grippers, since the gripper surfaces tend to mold slightly to the object surface. Single-point contact models also have a low tolerance to errors, especially errors of

the contact normal direction. Taking this into account a more realistic approach would be to employ a surface contact model which is not only more robust to errors but can also provide a better compliance with the object's surface geometry.

The first step is to preprocess the raw contact point cloud given by a physics engine in simulation or by a set of sensors on a real system. By simply calculating the mean of the raw contact point cloud and defining a single-point contact with the mean position, \mathbf{p}_m , and the mean normal direction this preprocessing is achieved in a simple and fast manner. To represent the mean normal direction, a new coordinate frame is created, R_m , with z axis along the mean normal direction and origin \mathbf{p}_m .

$$c_m = (\mathbf{p}_m, R_m) \quad (2.12)$$

Afterwards, a series of line segments parallel to the contact normal of c_m are generated. These line segments are bounded by a sphere centered in \mathbf{p}_m with radius r and by a plane containing \mathbf{p}_m and orthogonal to the contact normal of c_m . Calculating the intersections of each of the line segments with the object surface the points that define the contact surface are found. If multiple intersections are found for the same line segment the one closest to \mathbf{p}_m is chosen. The final step is to calculate the object surface normals of the points spanning the contact surface. The points in the contact surface are then modeled with the point-contact model described previously, each of them generating a set of wrenches $w_{c_i}^k$. Fig. 2.4 depicts the procedure described above.

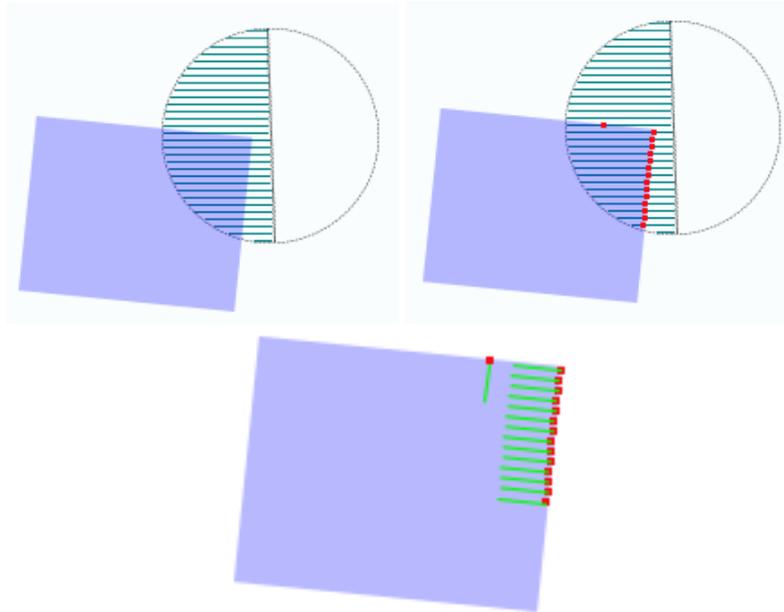


Figure 2.4: Steps for generating the Surface Contact Model. Line segments generated (top left), Line segments and their intersections with the object surface (top right) and intersection points and their corresponding surface normals (bottom). The blue box represents the object, the red points represent the intersections between the line segments and the object and the light green lines represent the object surface normal at each intersection point.

2.4 Grasp Representation

With the robust contact model framework introduced, defining a grasp is straightforward. Since all contacts are defined as sets of wrenches it is necessary to transform them to a common reference

frame R_r . The set of all the transformed wrenches defines the grasp, and is designated as grasp map G , [6]. Assuming that n contacts are generated

$$G_i = \begin{bmatrix} R_{c_i} & 0 \\ [p_{c_i}]_{\times} R_{c_i} & R_{c_i} \end{bmatrix} w_{c_i} \quad i \in [1, ..n]. \quad (2.13)$$

and the resulting grasp map is

$$G = [G_1, \dots, G_n]. \quad (2.14)$$

2.5 Closure

Force-Closure is a binary qualitative evaluation of grasp stability, [6]. If a grasp can resist any applied wrench it is considered force-closure. In other words given an external wrench $w_e \in \mathbb{R}^6$ applied to the object, there is a combination of contact forces f_c such that

$$Gf_c = -w_e \quad (2.15)$$

where

$$f_c = \begin{bmatrix} f_{c_1} \\ \vdots \\ f_{c_n} \end{bmatrix}, \quad f_{c_i} \in FC \quad (2.16)$$

and n is the number of contacts that compose the grasp. A simple way to evaluate if a grasp G is force-closure is given through the analysis of $ConvexHull(G)$. The $ConvexHull(G)$ represents the minimum convex region spanned by G on the wrench space W . Defining the grasp wrench space, W_G , as the space of all possible wrenches generated by the grasp

$$W_G = ConvexHull(G) \quad (2.17)$$

then G is force-closure if

$$W_0 \subset W_G \quad (2.18)$$

where W_0 is a small neighborhood of the wrench space origin.

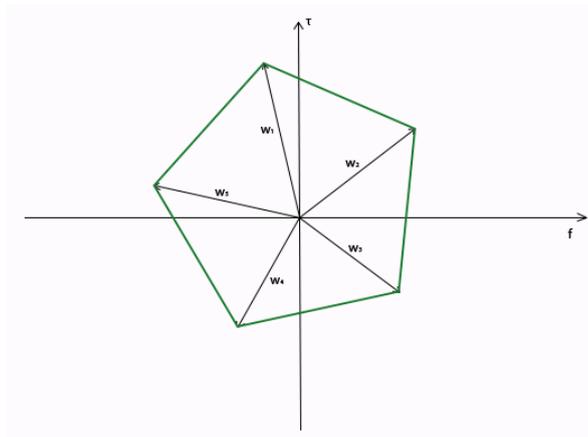


Figure 2.5: Convex Hull of the grasp wrenches, the Grasp Wrench Space

2.6 Bimodal Wrench Space Analysis

Typical metrics can give a measure of how good a force-closure grasp is but for non-force-closure grasps they give no information. Since for the set of all possible grasps that can be executed on a given object, only a small portion will pass the force-closure test, there is a large portion of the possible grasp set that remains unevaluated. The proposed metric attempts to solve this problem by evaluating all grasps even if they are non-force-closure, by means of different analysis modes chosen in accordance to the results of a force-closure test. Thus, for any configuration of the robot manipulator that touches the object (even if only with one finger) the metric will return a measure of grasp quality.

The proposed metric has two distinct analysis modes (one for force-closure and one for non-force-closure grasps) hence being called the Bimodal Wrench Space Analysis Metric (*BW*). The first mode, used for force-closure grasps, measures the radius of the largest sphere centered on the origin of the wrench space W , that is contained in W_G . Conceptually it represents the magnitude of the largest external force that can be resisted by the grasp in any direction. Because all the analysis is done under the assumption of unit normal force at all contact points, grasps where the sphere radius is larger have increased stability with the same amount of applied force.

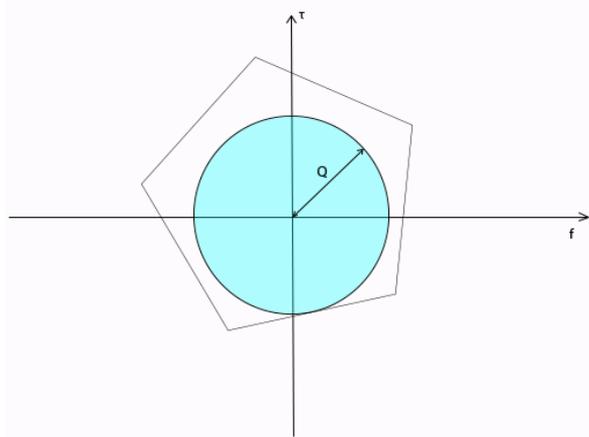


Figure 2.6: The Largest Sphere metric. Used by the *BW* metric to analyse force-closure grasps.

The second mode, used for non-force-closure grasps, measures the minimum distance from the origin of the wrench space W to a point contained in the W_G

$$\min_{\|y\|} : y \in W_G \quad (2.19)$$

Knowing that for a grasp to be force-closure the condition in (2.18) has to be met, it can be reasoned that, for non-force-closure grasps, if the grasp wrench space for grasp A, W_{G_A} , is closer to the origin than the grasp wrench space for grasp B, W_{G_B} , the changes in G_A in order to reach force-closure are inferior to the changes in G_B to reach the same condition. This is the reasoning behind the measure given by the second mode. In other words, mode two measures the distance to force-closure of a non-force-closure grasp.

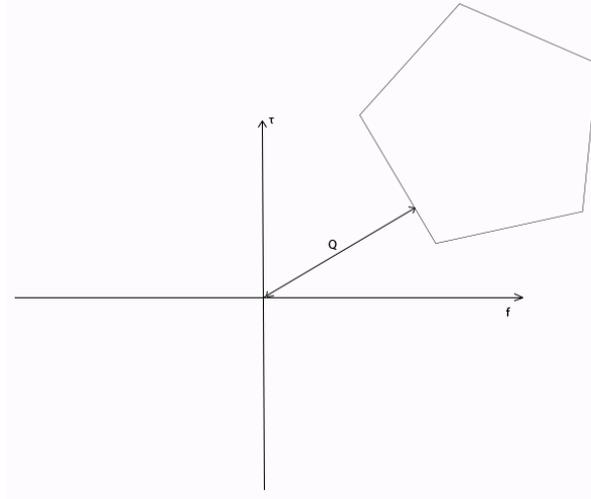


Figure 2.7: The BW metric. The grasp wrench space distance to force-closure when analysing a non-force-closure grasp.

A computationally interesting point of this second mode is that by representing the grasp wrench space region as

$$Nx \leq b \quad (2.20)$$

where N are the normal vectors and b the offsets that as a pair define each of the planes containing W_G , the second analysis mode can be represented as a simple convex optimization problem that can be solved easily by convex optimization solvers.

$$\min_{\|x\|} : Ax \leq b \quad (2.21)$$

To ensure smooth transitions between the two modes and hence a continuous metric function, the symmetric value given by mode two is used. This makes sense because as the distance to stability grows larger the overall evaluation of the grasp should be worse. On the other hand when the distance to stability tends to zero, the grasp is near transitioning to a force-closure and to positive metric values.

3

Bayesian Optimization

Contents

3.1	Gaussian Process Regression	18
3.2	Expected Improvement	19
3.3	Direct Optimization Algorithm	21
3.4	The Complete Method	22

Since the dawn of the computer age, one question that the scientific community strives to answer is "How can a computer learn?". With the emergence of robotics, machine learning arose as one of the most active scientific research areas.

Bayesian optimization methods were only recently introduced to machine learning applications, [3]. Their goal is to find the maximum value of an unknown function with as few function evaluations as possible. A series of complex stages are required to achieve this goal. The initial stages deal with predicting the completely unknown functions that describe the system while keeping track of the underlying uncertainty of these predictions. This is just the "tip of the iceberg". Later stages tackle how to learn efficiently and autonomously, improving the overall knowledge of the system. All of this coupled with dealing with the exploration versus exploitation trade-off and maintaining the computational feasibility of the method is not a trivial task.

In this chapter all of the stages that compose this method will be described in detail. Gaussian Process regression will give us an approximation of the grasp metric function and is the first topic to be discussed. The active learning by means of the expected improvement assessment will reduce the number of exploration function calls needed and is the second focus of this chapter. Finally the Direct Optimization algorithm is discussed as a means to improve the overall efficiency of a method that without it suffers from the Curse of Dimensionality. Finally, all things together form an efficient optimization method for generic unknown functions.

3.1 Gaussian Process Regression

When trying to learn one always needs something to learn from. This something must provide some sort of new information about the subject of interest in a way that the learning agent understands. In the case of a robot trying to learn an unknown function that represents the problem it is facing, it will clearly need some sort of function representation in order to learn. This is what a Gaussian process provides.

A Gaussian is used as a means to describe a distribution over functions. As a more formal definition, a Gaussian Process (*GP*) is a collection of random variables, any finite number of which have a joint Gaussian distribution, [10]. It is completely defined by a mean and covariance function pair where

$$\mu(x) = \mathbb{E}[f(x)] \tag{3.1}$$

$$\Sigma(x, x') = \mathbb{E}[(f(x) - \mu(x))(f(x') - \mu(x')))] \tag{3.2}$$

and is represented as

$$f(x) \sim \mathcal{GP}(\mu(x), \Sigma(x, x')). \tag{3.3}$$

In the current context the random values represent the function values $f(x)$ that the robot wishes to learn. The mean function is the best prediction of the true function given what is known. Also it is initially considered as a zero function since there is no relevant information at the start of the process, but other priors may be used if desired.

For representing the covariance function several kernels may be used. A kernel is the general name given to a function K of two arguments mapping a pair of inputs $x \in X, x' \in X$ into \mathbb{R} . In this work the kernel chosen is a Matérn class covariance function given by

$$K(r) = \left(1 + \frac{\sqrt{3}r}{l}\right) \exp\left(-\frac{\sqrt{3}r}{l}\right) \quad (3.4)$$

where

$$r = |x - x'| \quad (3.5)$$

measures the distance between points x and x' and l represents the kernel length. This length is the minimum distance that one has to travel along the input arguments for the variation of the function value to be noticeable. For more information on this kernel refer to [10]. The representation of a covariance function is what implies a distribution over functions. Hence at the beginning of the *GP* one can predict several randomly generated functions that are possible through

$$\mathbf{F}_* \sim \mathcal{N}(0, K(\mathbf{X}_*, \mathbf{X}_*)) \quad (3.6)$$

where \mathbf{X}_* is the estimation point set, \mathbf{F}_* denotes the corresponding estimated function values and $K(\mathbf{X}_*, \mathbf{X}_*)$ denotes the kernel evaluated between all point pairs of the prediction point set.

Although possible, these estimations are not very interesting since there is no real information taken into account thus spanning infinite possible solutions. To consider points observed from $f(x)$ when performing the estimation will clearly provide a better insight on the function. To do this it is a simple matter of conditioning the distribution to what is already known

$$\mathbf{F}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{f} \sim \mathcal{N}(K(\mathbf{X}_*, \mathbf{X})K(\mathbf{X}, \mathbf{X})^{-1}\mathbf{F}, K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X})K(\mathbf{X}, \mathbf{X})^{-1}K(\mathbf{X}, \mathbf{X}_*)). \quad (3.7)$$

where \mathbf{X} is the observation (known) point set and \mathbf{F} is the set of observed function values. The values $K(\mathbf{X}_*, \mathbf{X})$, $K(\mathbf{X}, \mathbf{X}_*)$, $K(\mathbf{X}, \mathbf{X})$ are the kernel evaluated between point pairs of the prediction and observation set.

Despite the fact that this is a fairly good estimation of the function it is still somewhat naive. The thought of getting perfect measurements experimentally is an extremely gullible approach. Therefore some observational error has to be taken into account when performing the estimation. Assuming additive independent identically distributed Gaussian noise ε with variance σ_n^2 leads to

$$\mathbf{F}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{F} \sim \mathcal{N}(K(\mathbf{X}_*, \mathbf{X})(K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I)^{-1}\mathbf{F}, K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X})(K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I)^{-1}K(\mathbf{X}, \mathbf{X}_*)) \quad (3.8)$$

where I is the identity matrix.

This estimation of the mean function provides us with the robust estimation of the real function that we need. It also provides a measure of the estimation's uncertainty, which as we will see in the next section will be crucial to the robot's learning.

3.2 Expected Improvement

Now that we have something that provides a solid representation of what we wish to learn and that allows us to predict what is still unknown based on current knowledge, the next step is to decide

what we should do to better improve this knowledge. This decision shouldn't be taken lightly since it is what defines the rate at which we learn. To achieve this decision we will use the concept of Expected Improvement [3] that will be introduced in this chapter.

The Gaussian Process provides a global estimation of $f(x)$ based on what is known at the time. Since we are trying to find the maximum value of the function $f(x)$ the natural decision should be to explore regions of the function where a higher maximum is possible. From this idea we define the improvement function has

$$I(x) = \max\{0, f_{n+1}(x) - f^{max}\} \quad (3.9)$$

where f^{max} is the current maximum value. The function takes on positive values when the prediction is higher than the best value found so far and is set to zero otherwise. Using the improvement function, the new observation point is attained by finding the maximum expected improvement point

$$x = \arg \max_x \mathbb{E}(\max\{0, f_{n+1}(x) - f^{max}\} | \mathcal{D}_n) \quad (3.10)$$

where \mathcal{D}_n is all that is known at time n

$$\mathcal{D}_n = [f_1, \dots, f_n] \quad (3.11)$$

and f_i is the observed value for trial i . This expected improvement can easily be evaluated analytically, [3], through

$$EI(x) = \begin{cases} (\mu(x) - f^{max})\Phi(Z) + \sqrt{\text{diag}(\Sigma)}\phi(Z) & \text{if } \text{diag}(\Sigma) > 0 \\ 0 & \text{if } \text{diag}(\Sigma) = 0 \end{cases} \quad (3.12)$$

where

$$Z = \begin{cases} \frac{\mu(x) - f^{max}}{\sqrt{\text{diag}(\Sigma)}} & \text{if } \text{diag}(\Sigma) > 0 \\ 0 & \text{if } \text{diag}(\Sigma) = 0 \end{cases}, \quad (3.13)$$

Σ is the covariance function and $\phi(\cdot)$ and $\Phi(\cdot)$ respectively denote the probability density function and the cumulative distribution function of the standard Normal distribution. As mentioned in the previous section, the uncertainty measure given by the covariance function K plays a huge role on the decision of the next observation. It enables the balancing between exploiting and exploring. When exploring, we should focus on points where the prediction variance is large in order to minimize the global uncertainty. When exploiting one's focus should be the points where the predicted mean function is high so that a higher and more accurate value of the global maximum may be found.

Evaluating the expected improvement through (3.12) balances the exploration versus exploitation trade-off in an unruly fashion. A more generalized form of the $EI(\cdot)$ has to be found in order to directly control this balance. We arrive at

$$EI_\xi(x) = \begin{cases} (\mu(x) - (f^{max} + \xi))\Phi(Z_\xi) + \sqrt{\text{diag}(\Sigma)}\phi(Z_\xi) & \text{if } \text{diag}(\Sigma) > 0 \\ 0 & \text{if } \text{diag}(\Sigma) = 0 \end{cases} \quad (3.14)$$

where $\xi \geq 0$ and

$$Z_\xi = \begin{cases} \frac{\mu(x) - (f^{max} + \xi)}{\sqrt{\text{diag}(\Sigma)}} & \text{if } \text{diag}(\Sigma) > 0 \\ 0 & \text{if } \text{diag}(\Sigma) = 0 \end{cases}. \quad (3.15)$$

In this work we use $\xi = \frac{\hat{\sigma}_f^2}{100}$ where $\hat{\sigma}_f^2$ is the Gaussian process estimated variance given by

$$\hat{\sigma}_f^2 = \mathbf{F}^T K(\mathbf{X}, \mathbf{X})^{-1} \mathbf{F}. \quad (3.16)$$

For a more detailed reading on this topic refer to [3].

As it now stands, the method seems complete. We have a procedure to represent the current knowledge, to predict the unknown and to accordingly decide what to do next. Repeating this process will ultimately lead to learning $f(x)$ very efficiently in terms of minimizing the number of observations. The last section of this chapter will reveal the last problem to be tackled by the method.

3.3 Direct Optimization Algorithm

Now that the method is fairly complete, one question arises. Is it computational feasible? The answer is yes, with complexity $O(n)$, where n is the number of points estimated. Despite the low complexity of the system, it struggles for higher dimensional inputs. This stems from the fact that to calculate the Expected improvement one needs to calculate several kernels. These kernels must be evaluated at all point pairs possible between the points in the estimated and observation point sets. While the kernel evaluation is simple, the number of points in the estimated point set grows at troubling rates when doing uniform sampling through space. This point set grows at a rate of n^D where n is the number of samples per dimension and D is the number of input dimensions. And so if working with 2-dimensional inputs and estimating 10 points per dimension, 10^2 points are generated. On the other hand the same 10 points per dimension for a 6-dimensional inputs space span 10^6 points. This is commonly known as the "Curse of Dimensionality". This means that simple uniform sampling of the space along all the dimensions in order to find the point with the highest expected improvement is not a good approach.

We turn to a more efficient approach through the use of the Direct Optimization algorithm, [11]. This algorithm uses a small number of initial predictions to decide how to divide the feasible space into smaller RECTangles. The end result is a high discretization of the target function near the function maxima and a low discretization elsewhere.

The Direct algorithm starts by normalizing the function domain into a unit hyper-cube with center c_1

$$\bar{\Omega} = \{x \in \mathbb{R}^N : 0 \leq x \leq 1\}. \quad (3.17)$$

The algorithm works in this normalized space, only reverting to the original space when making function calls. After evaluating the function at $f(c_1)$ it is time to make the first division of the hyper-cube. The cube is divided into smaller cubes centered at $c_1 \pm \delta e_i$, $i = 1, \dots, n$ where δ is one third of the cube length and e_i is the i th unit vector. Direct chooses to leave the best function values in the largest space. As such the first dimension to be divided is chosen by means of

$$\omega_i = \min(f(c_i + \delta e_i), f(c_i - \delta e_i)), \quad 1 \leq i \leq N. \quad (3.18)$$

The dimension with the smallest ω_i is divided into thirds and the process is repeated for all dimensions on the resulting center hyper-rectangles. Evaluating the function at all the resulting center points, $f(c_i)$ brings the algorithms initialization to a close.

With the hyper-cube division done it is time to find which of the newly generated rectangles/cubes may be potentially optimal. For the optimality test, we test each of the rectangles/cubes for the

existence of a Lipschitz Constant $\hat{K} > 0$ that allows

$$f(c_j) - \hat{K}d_j \leq f(c_i) - \hat{K}d_i, \forall i, \quad (3.19)$$

$$f(c_j) - \hat{K}d_j \leq f_{min} - \epsilon|f_{min}| \quad (3.20)$$

where $\epsilon > 0$ is a positive constant, f_{min} is the current best function value and c_j and d_j are respectively the center of the tested hyper-rectangle/cube and a measure of the dimension of the same rectangle/cube. In (3.19) we test if the possible variation of the value $f(c_j)$ when traveling inside the respective hyper-rectangle/cube may reach a minimum value when applying the same variation scale to all other $f(c_i)$ navigating on the respective i th rectangle/cube. On the other hand (3.20) test if the possible minimum reached on the j th rectangle/cube is of interest when compared with f_{min} . The term $\epsilon|f_{min}|$ makes sure that the improvement to the minimum value is non trivial.

Now that we know S , the set of all the potentially optimal rectangles/cubes, the only remaining step is to divide all members of this set, evaluate the function at the center of the resulting rectangles/cubes and update f_{min} . An interesting fact is that when dividing an hyper-rectangle, the algorithm always chooses to divide along the longest dimension(s) to ensure that the rectangles shrink on every dimension.

The Direct algorithm repeats the previous operations (except the initialization) until no S is empty, meaning no more divisions are of interest. When this stage is reached the f_{min} is the global function minimum. Since we are looking to maximize the expected improvement function, it is only a matter of suppling Direct with negative values of $EI(\cdot)$.

Having presented all the tools involved, the next section shows how they are all integrated with each other to make the complete method.

3.4 The Complete Method

Now that all the tools that compose the method have been presented, it remains to show the sequence in which they are used to bring about the desired results. So what steps are necessary in order for the method to preform an intelligent mapping of robot manipulator parameters to Grasp quality metrics, while ensuring that the mapping captures the maximum values reachable in the available parameter space?

The first step is to acquire a sample of the mapping in order to initialize the gaussian process. Since there is no prior information, the input parameters for this sample are chosen randomly.

Having the gaussian process initialized we start searching for the next best point to sample using the expected improvement. Here is where the Direct comes into play. The Direct algorithm will provide the gaussian process with the points that it has to predict in order to reach the maximum value of the expected improvement function and the resulting parameters that describe the next sample point.

Note that Direct does not supply the Gaussian process with all prediction points at once. The points will be provided progressively as the algorithm divides the parameter space. As the points are given to the Gaussian process, the expected improvement is calculated immediately so that Direct may continue working.

Knowing the next point to be sampled, the robot is called to evaluate it and the new observation is introduced to the Gaussian process. From here the cycle restarts and starts searching for the next sample.

This cycle ends when the expected improvement is 0, meaning there are no points that will provide additional relevant information to the system.

Figure 3.1 shows the flowchart of the method depicting what was described above.

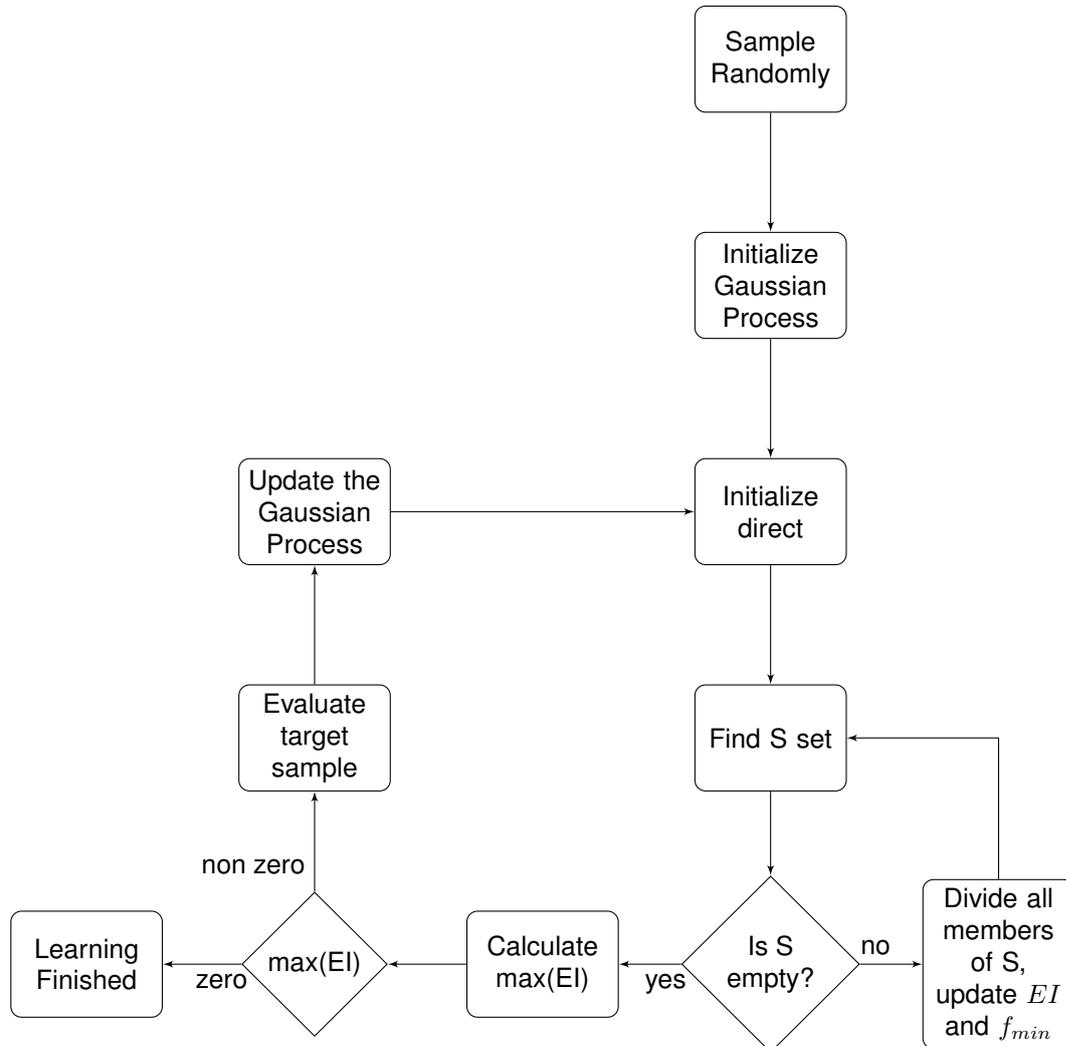


Figure 3.1: Flowchart depicting the complete Bayesian Optimization method.

4

Experimental Results

Contents

4.1	Experimental Setup	26
4.2	Metric Comparisons	27
4.3	Contact Model Influence	29
4.4	Reference Frame Independence	30
4.5	Searching with two parameters	31
4.6	1D exploration with Bayesian Optimization	34
4.7	2D exploration with Bayesian Optimization	36
4.8	Bayesian optimization versus random sampling	40
4.9	Results on a real platform	42

In this chapter the experimental results obtained throughout this work will be presented. The main purpose of the experiments in a first stage was to evaluate the proposed metric, Bimodal Wrench Space Analysis (BW). This was done by placing the hand on a reference position close to the object, closing the fingers until contact with the object surface, and computing the several grasp quality metrics for comparison purposes. Comparison and evaluation of the grasp metrics was done empirically through the observation of $1D$ and $2D$ plots of their values. Also, some of the important aspects of the assessment process (the reference frame independence and the surface model influence) will be demonstrated through the experiments. Closing the first stage of experiments, the behavior of the BW metric is shown for several objects. This is achieved by changing the reference position of the hand with respect to the object in small steps along the approach, x and vertical y directions in order to attain a dense $2D$ sampling of the results, Fig. 4.1

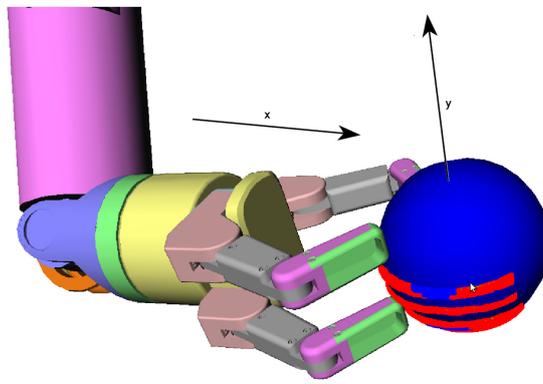


Figure 4.1: 2D sampling of a Sphere

The second stage of the experiments pursues the testing of the Bayesian Optimization method's performance. A comparison is also made between Bayesian optimization methods and a simpler method. To illustrate the process used by the Bayesian optimization method we start by sampling the proposed metric function in $1D$ space. We then proceed to test the sampling efficiency of the method when the input space is 2 -dimensional. We finish our analysis of the Bayesian Optimization method by comparing how quickly it converges to the optimal functions values when compared to simple random sampling through the input space.

The last experiment shows preliminary results of integrating this work on the HANDLE project platform.

4.1 Experimental Setup

Most of the experiments performed during the course of this work were done in a simulation environment through the OpenRave simulator [2]. The environment consisted of a manipulator arm, the Barrett hand, and a few objects as shown in Fig. 4.2. The calculation of convex hull's, required for the calculation of the proposed metric, required the use of the qhull library. The free convex optimization library for the python language, cvxopt, was the tool chosen to solve the convex optimization problem that leads to the second mode of the proposed metric.

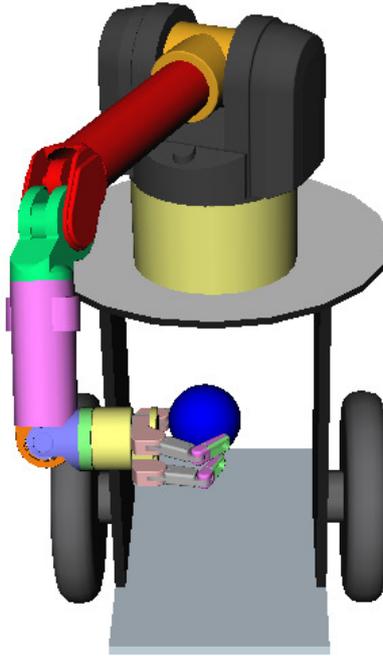


Figure 4.2: The experimental setup used for validating the proposed metric.

4.2 Metric Comparisons

In this section two well known metrics, Largest Sphere (LS) [8] and Grasp Wrench Volume (GWV) [7], are compared with the proposed Bimodal Wrench metric (BW). Comparisons are based on 1D plots obtained by sampling the metrics along the hand approach direction to the object, x . The first test, shown in Fig. 4.3, compares BW , LS and GWV , as the hand scans a sphere (the sphere object is the same as in Fig. 4.9).

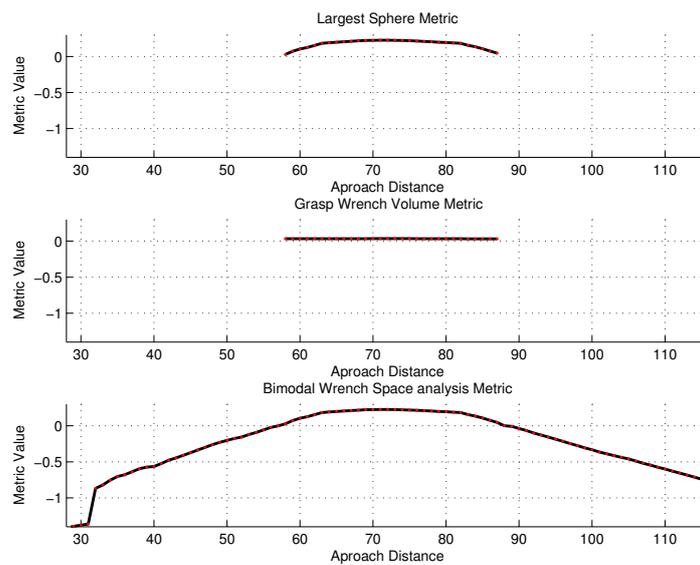


Figure 4.3: Comparison between Largest Sphere Metric (top), Grasp Wrench Volume Metric (middle) and BW Metric (bottom) for grasps on a sphere.

Being a smooth and very regular object, both LS and BW exhibit a smooth behavior with a clear maximum at about the center of the object while the GWV , despite clearly identifying the force-closure region, does not provide any additional information. Note that whereas the classical metrics only provide results in the force closure range (about $30mm$), our metric provides sensible results in a range of more than $80mm$. It is visible a smooth transition from positive (force-closure) to negative (non-force-closure) values, which provide continuity in the transition of the employed semi-metrics.

The following test performs the same type of analysis on a $45deg$ rotated cuboid (this object is also shown in Fig.4.14). Results are shown in (Fig. 4.4). While the hand approaches the object it first starts contacting the closest edge, then gradually the contact points are set apart until the fingers reach the two lateral edges. During a span of about $20mm$ force-closure grasps are attained. After this point, the fingers pass to the other side of the middle edges. In this case, despite there being a discontinuity in the BW metric the relative assessment between the force-closure and non-force-closure segments is possible. Also, seeing the cuboid's edges has a discontinuity in the metric may be useful in some applications.

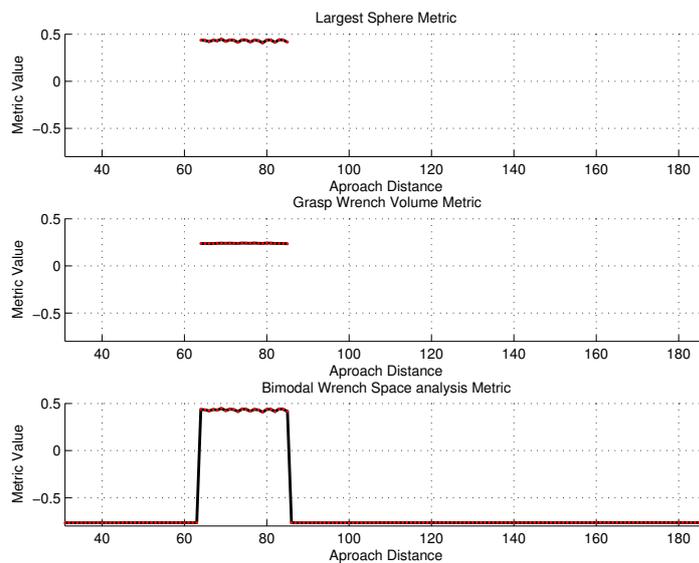


Figure 4.4: Comparison between Largest Sphere Metric (top), Grasp Wrench Volume Metric (middle) and BW Metric (bottom) for grasps on a rotated cuboid.

The last test on this set compares the metric results when performing a scan of a star prism (this object is the same as shown in Fig. 4.15). The results shown in Fig. 4.5 suggest that discontinuities in the metric can actually be viewed as abrupt changes in the objects shape.

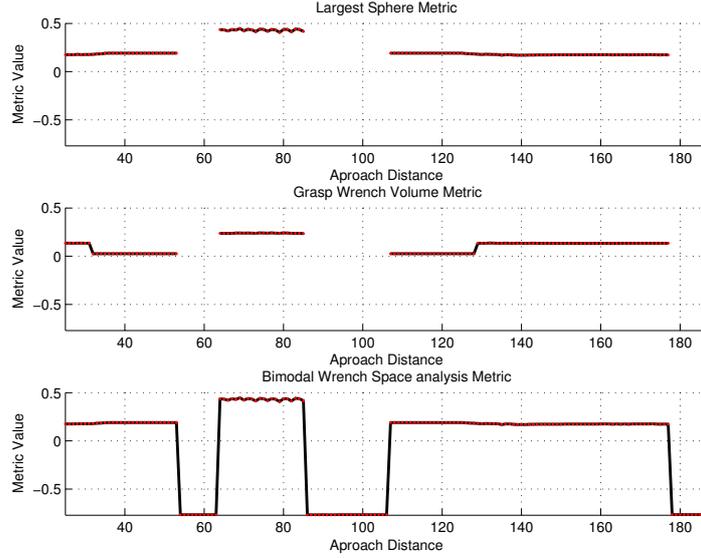


Figure 4.5: Comparison between Largest Sphere Metric (top), Grasp Wrench Volume Metric (middle) and BW Metric (bottom) for grasps on a star prism.

4.3 Contact Model Influence

In this section we compare the results of the proposed metric when used with different contact models. The first test analyses the variation of the metric results when using four different types of contact models. The two first models differ as result of the mean preprocessing done on the raw point cloud contacts. The third model uses contact points on a circular plane containing the mean position p_m and orthogonal to the mean normal. The last model uses the full surface model described on Sec. 2.3. Results are shown in Fig. 4.6.

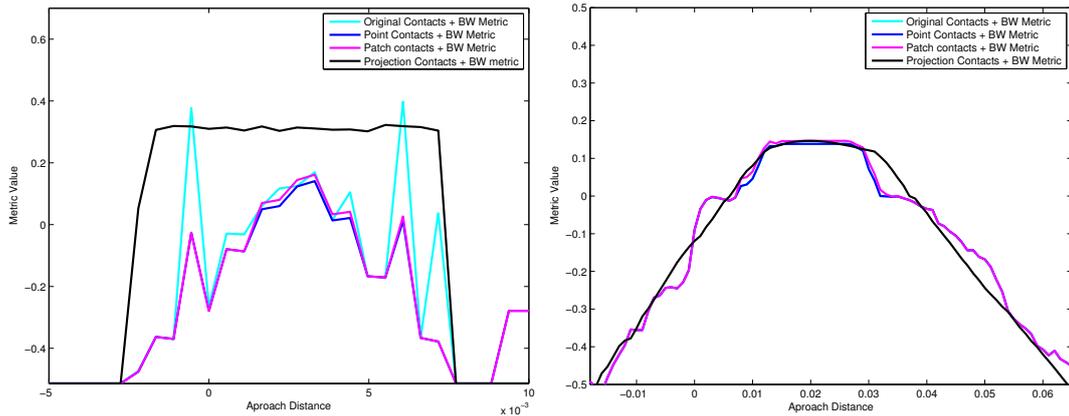


Figure 4.6: BW Metric values using different contact models when executing grasps on a cuboid (left) and on a sphere (right).

It is clear that the more complex the contact point representation is, the better the results are. For both objects the results obtained when using the full model undoubtedly give a more realistic representation of the grasp's variation along the object. In Fig. 4.7 the improvements are further emphasized by making a 2D scan of a sphere while using only single-point contacts and the full

surface model.

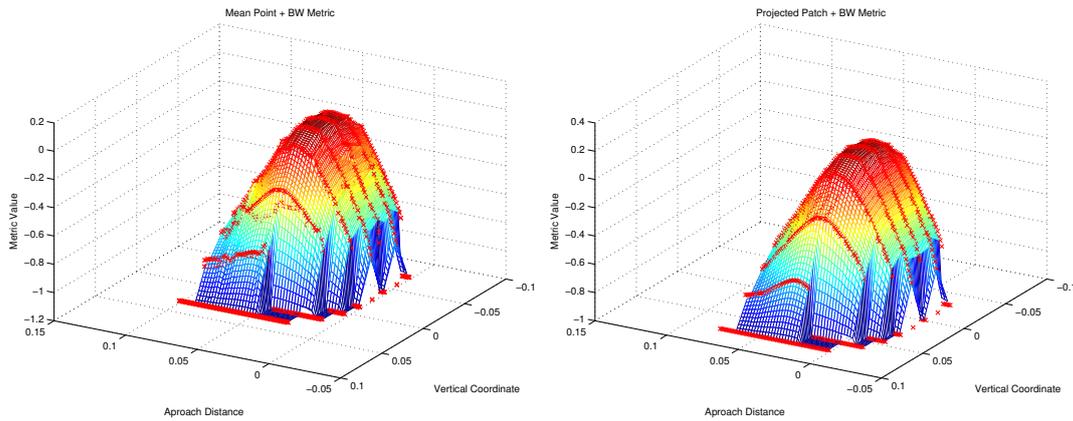


Figure 4.7: *BW* Metric values of grasps in a sphere, computed in a grid sampling of the hand position parameters with the mean point model (left) and with the projected patch model (right).

4.4 Reference Frame Independence

This section shows the effects of changing the reference frame for contact point representation. As shown in Fig. 4.8, the shape of the metric curves obtained and the position of the maximum value will generate the same grasp hierarchy independently of the changes to the reference frames. The important aspect of this analysis is the following. When trying to compare a set of grasps using this representation, as long as the reference frame chosen at the beginning of the experiment remains unchanged throughout the experiment's duration, any frame can be chosen as reference. This allows this system to operate with no prior knowledge of the object it is trying to grasp.

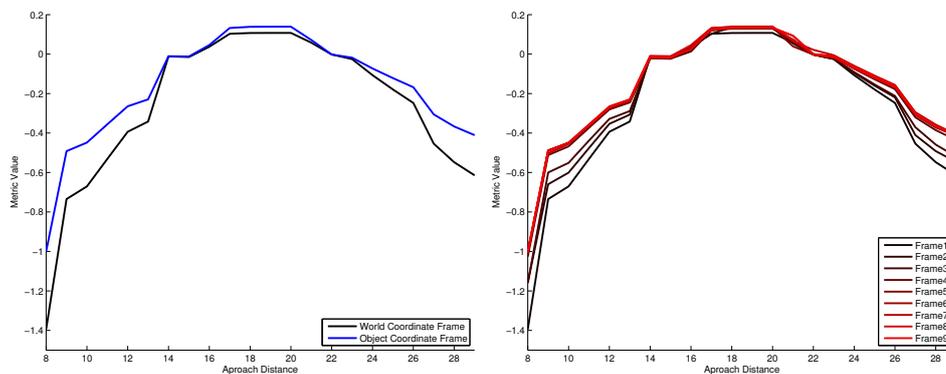


Figure 4.8: Metric Results using as reference frame the World frame and the Object frame (left) and some of the robot's frames (right)

This property allows the use of this work on non-model based approaches. Coupled with some other results to be shown in the next section, the use of this system can give some insight on the unknown object's basic shape.

4.5 Searching with two parameters

In this section a series of 2D “scans” are made for different objects. It is shown that whenever dealing with smooth surfaces (figures 4.9 to 4.12) the scans depict the objects through fluid continuous transitions.

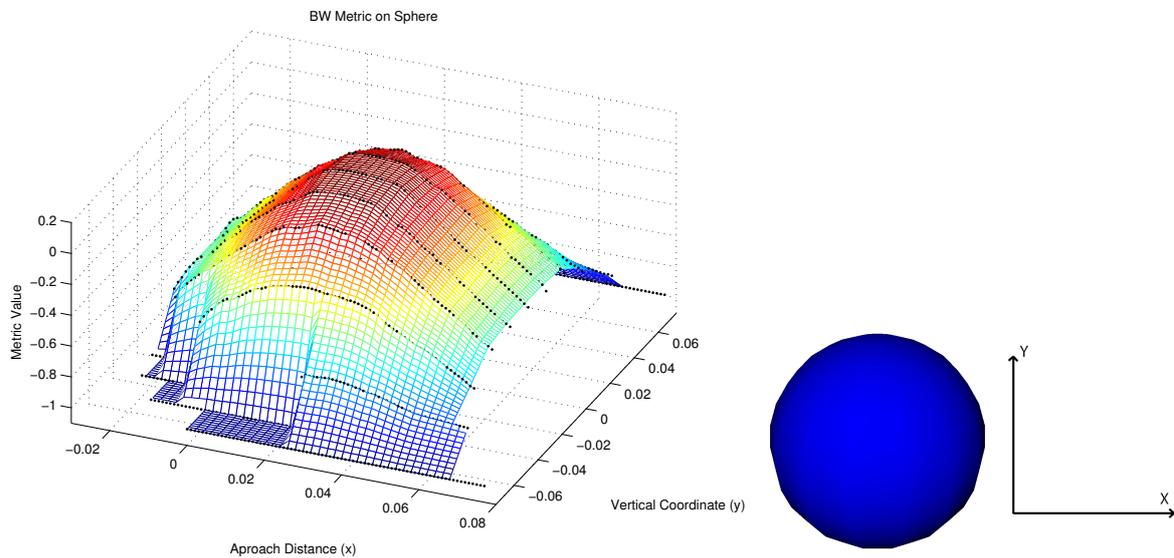


Figure 4.9: *BW* Metric values of grasps in a sphere, computed in a grid sampling of the hand position parameters.

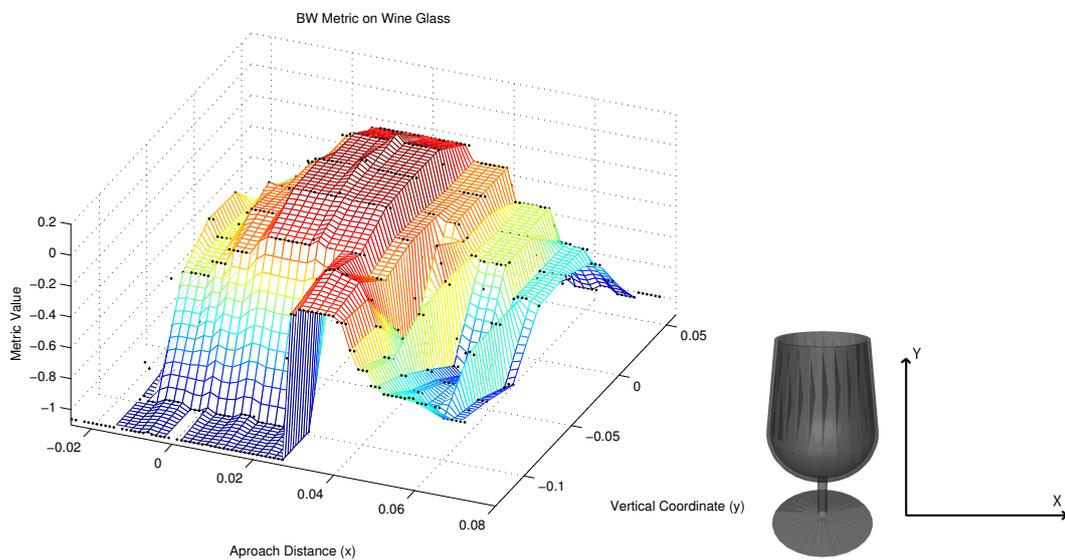


Figure 4.10: *BW* Metric values of grasps in a wine glass, computed in a grid sampling of the hand position parameters.

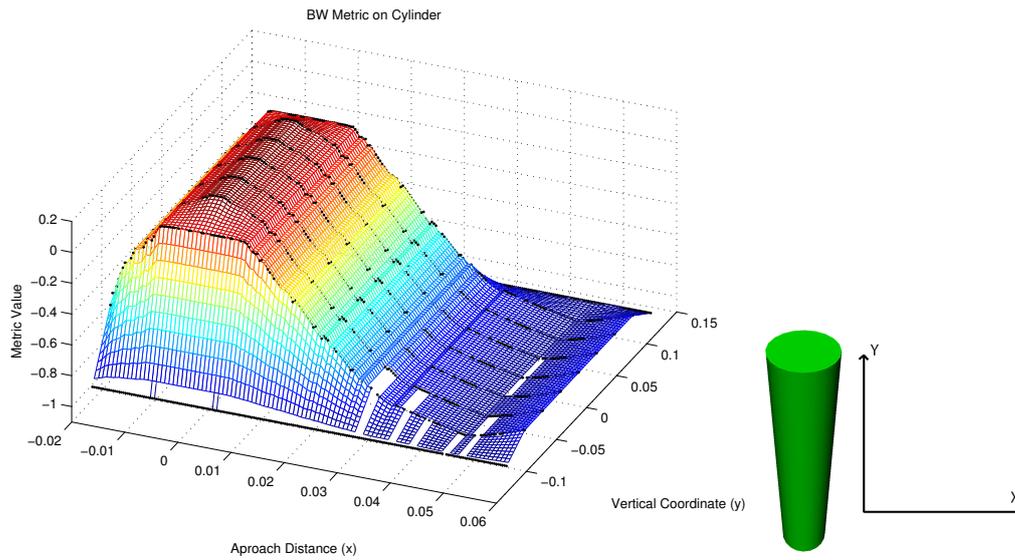


Figure 4.11: *BW* Metric values of grasps in a cylinder, computed in a grid sampling of the hand position parameters.

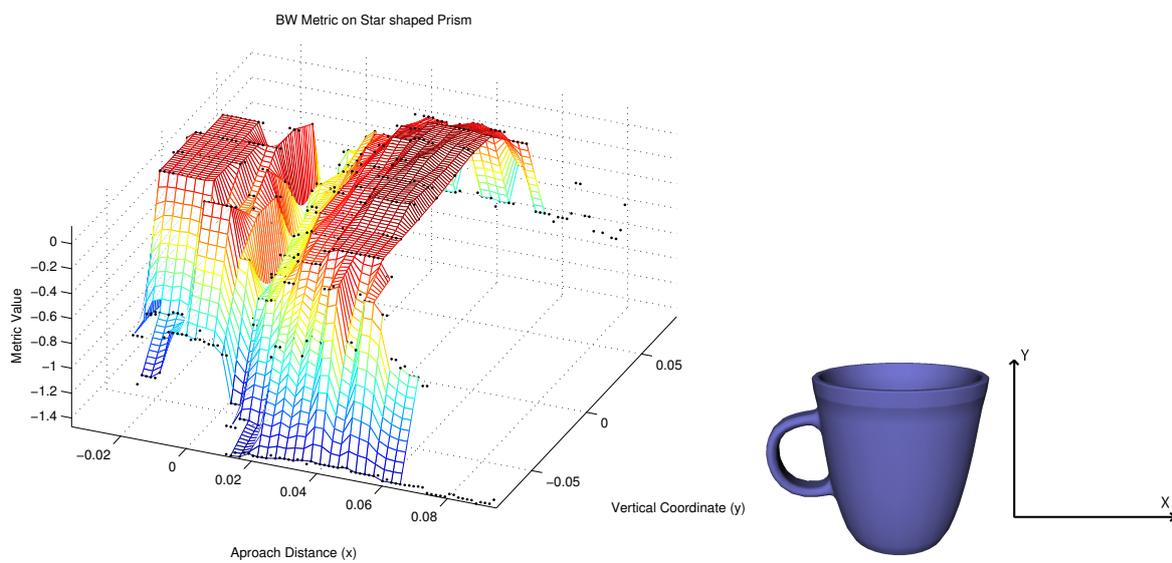


Figure 4.12: *BW* Metric values of grasps in a mug, computed in a grid sampling of the hand position parameters.

On the other hand when the objects exhibit sharp edges (figures 4.13 to 4.15) the scans show hard discontinuities in the edge regions and fluid transitions on the rest of the object.

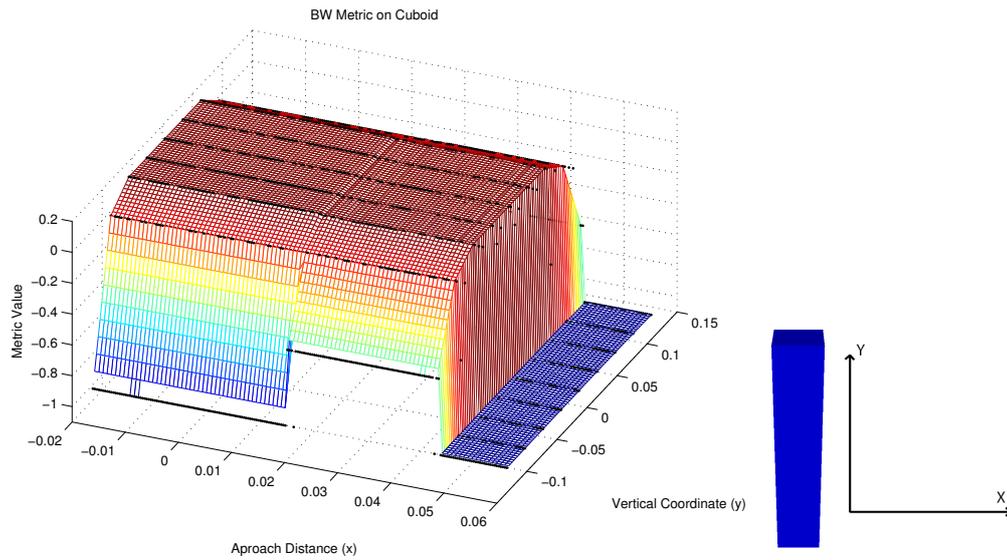


Figure 4.13: *BW* Metric values of grasps in a cuboid, computed in a grid sampling of the hand position parameters.

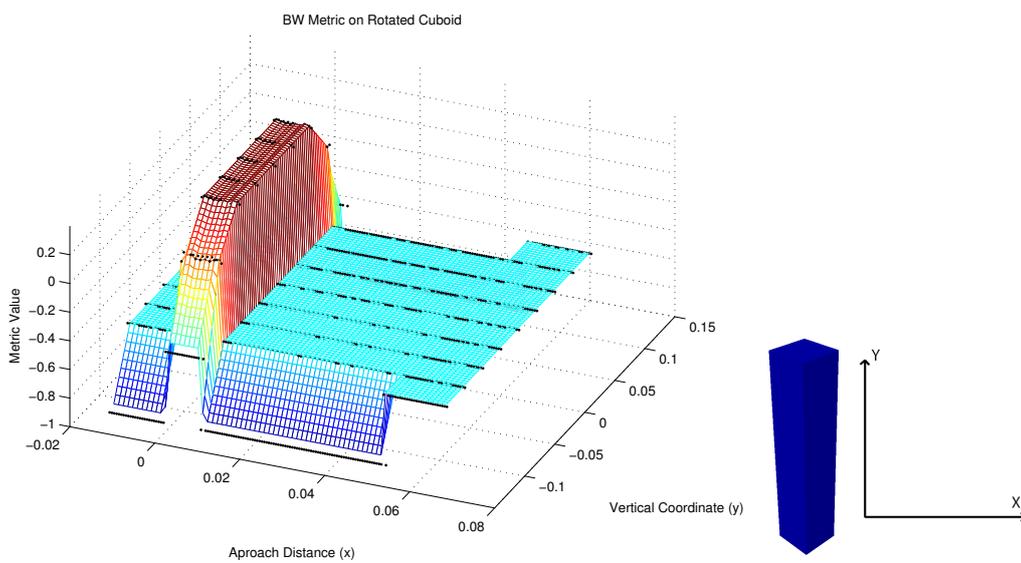


Figure 4.14: *BW* Metric values of grasps in a rotated cuboid, computed in a grid sampling of the hand position parameters.

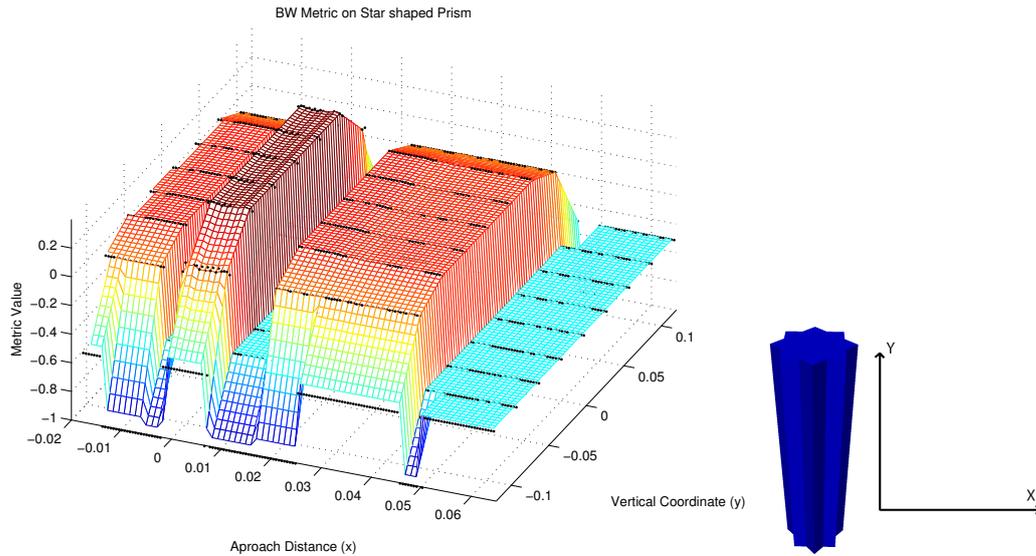


Figure 4.15: *BW* Metric values of grasps in a star prism, computed in a grid sampling of the hand position parameters.

Thus, it is shown that the metric provides a good grasp characterization for both force-closure and non-force closure grasps independently of the object shape. Additionally by solely analyzing the metric's outputs, some basic inference on the object's shape can be made.

4.6 1D exploration with Bayesian Optimization

In this section we show some 1D scans made by changing the initial position along the approach axis. The goal of these scans is to give a better understanding of how the Bayesian methods work. Also to be shown is how the system behaves when the initial random sample is one of the best possible or one of the worst possible. Fig. 4.16 depicts the latter. It represents the sequence carried out by the method to sample a sphere, when the initial random sample does not even touch the object (top left plot). In red we can see the *GP* mean function which is the best estimation of the metric function at this time, the dashed lines depict the estimation variance at each function point, the red dots are the values collected from the robot trials. The blue function represents the *EI* function that classifies the function space in terms of exploration interest. Even with the setback caused by the bad initial random sample, it only takes 2 iterations of the method in order to find a possible maximum (top right plot). Acknowledging the fact that a region that may contain the maximum has been found, the system focuses the search in it's neighboring points (bottom left plot). Once this region has been exploited, the system resumes it's exploration efforts to assure there are no more regions that may contain better values of $f(x)$. After confirming that it has found the function maximum value the exploration stops (bottom right plot).

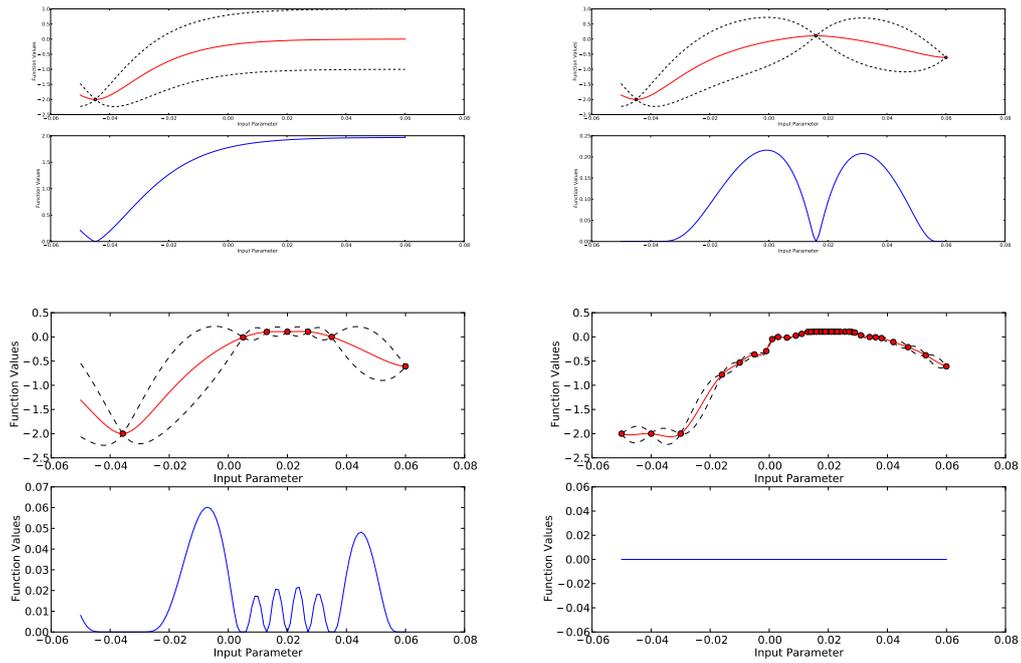


Figure 4.16: 1D scan of a sphere, computed by sampling along the input parameter.

The second case that will be shown is the opposite case. Fig. 4.17 shows the sequence followed by the method while sampling a star prism when the initial random sample is one possible maximum (top left plot). As the function is still completely unknown to the system, excluding the random sample, it is impossible for it to realize that the first sample is actually a possible maximum and so it proceeds with the exploration. After 3 iterations the system finally realizes the potential of the first sample (top right plot) examining the neighboring region (bottom left plot). The exploration is resumed and 2 more interest areas are found before the systems stops (bottom right plot).

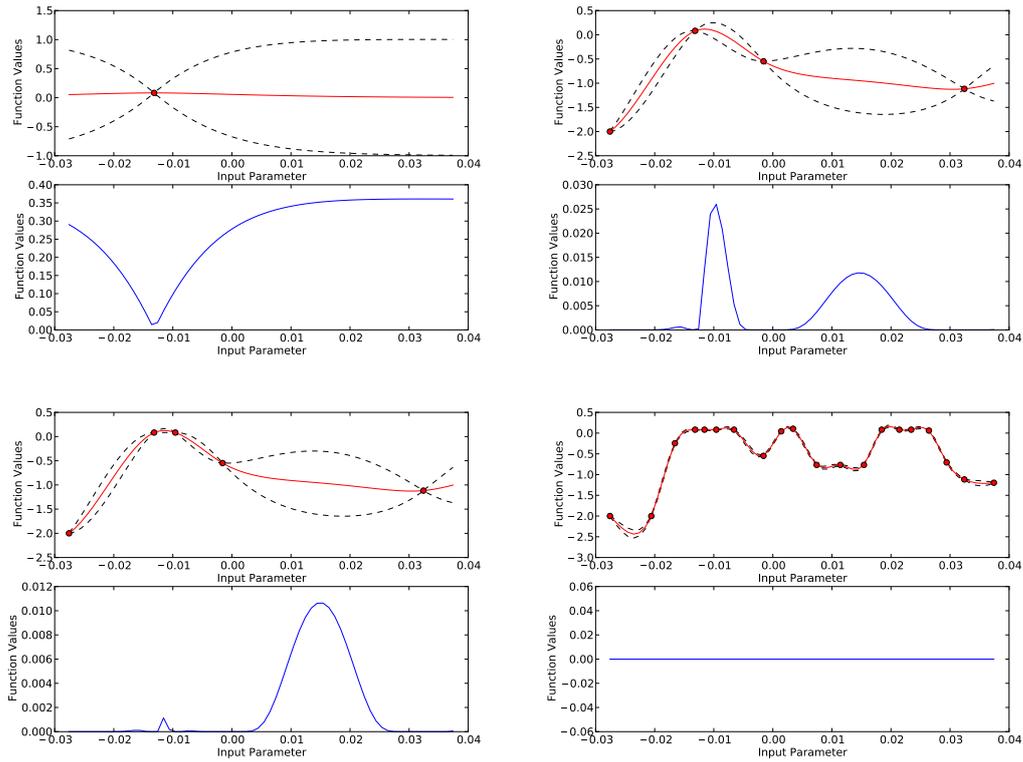


Figure 4.17: 1D scan of a star prism, computed by sampling along the input parameter.

BW

4.7 2D exploration with Bayesian Optimization

In this section we will show the results of scanning the object in the same 2-dimensional space as in section 4.5 but using the Bayesian optimization method. Also in order to test the method performance under the worst possible conditions, we assume only point contacts are available and that the observations are noisy.

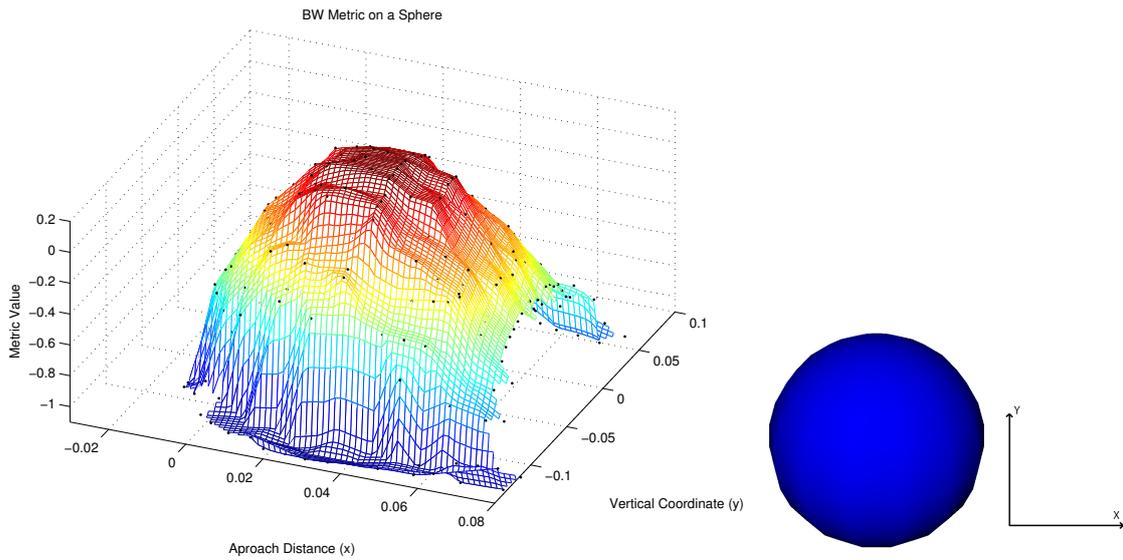


Figure 4.18: *BW* Metric values of grasps in a sphere, computed using the Bayesian Optimization method.

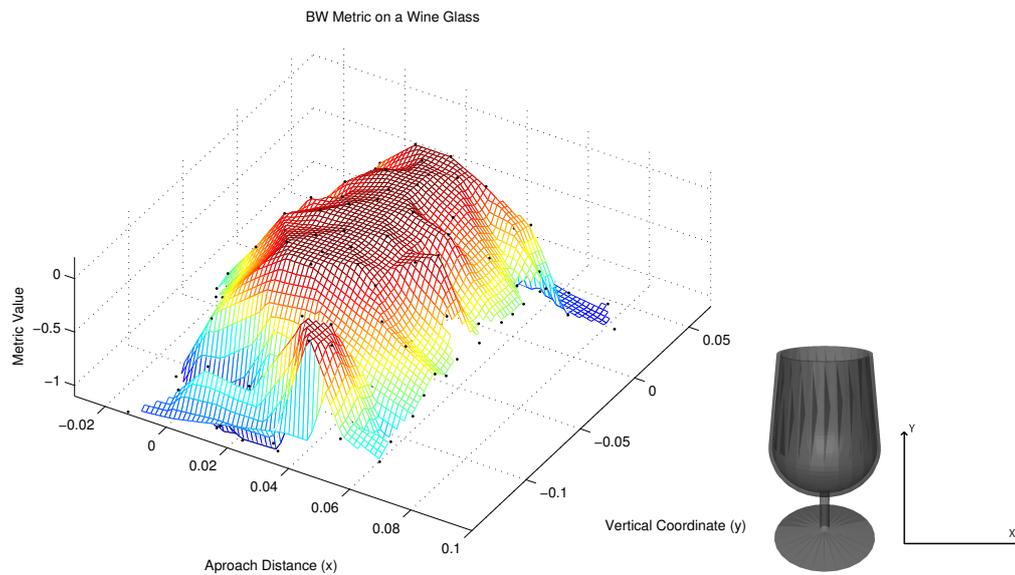


Figure 4.19: *BW* Metric values of grasps in a wine glass, computed using the Bayesian Optimization method.

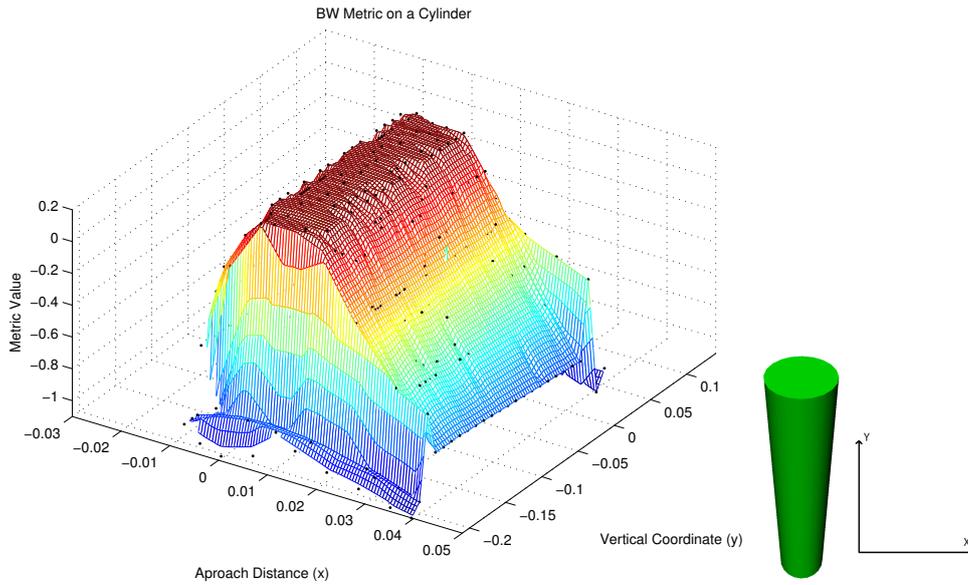


Figure 4.20: *BW* Metric values of grasps in a cylinder, computed using the Bayesian Optimization method.

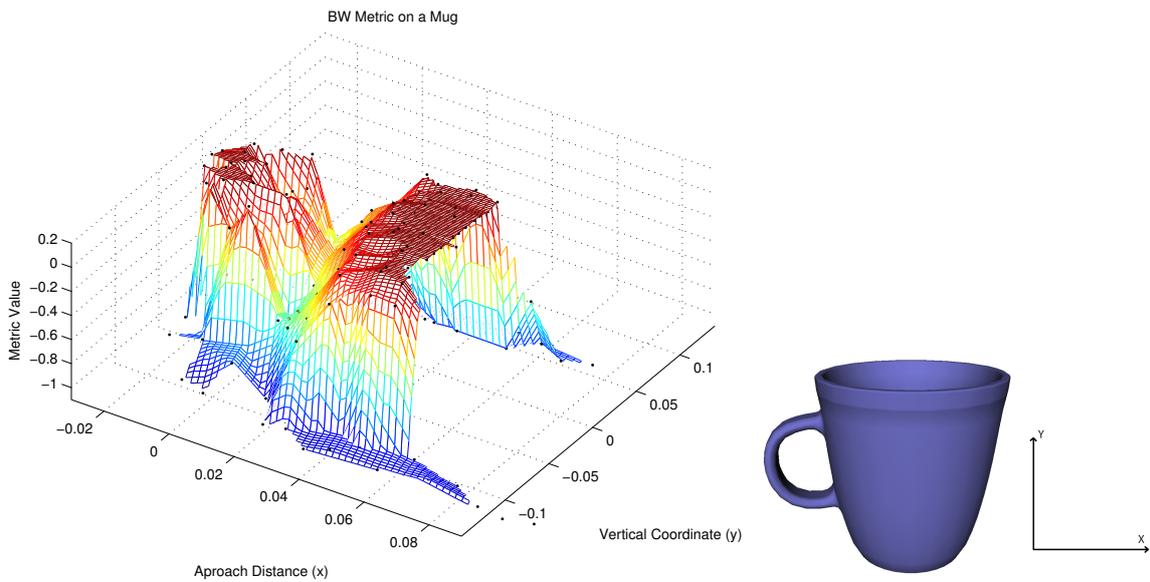


Figure 4.21: *BW* Metric values of grasps in a mug, computed using the Bayesian Optimization method.

We can observe all the metric properties that were referred in section 4.5 although the presence of an edge is now translated to a smoother transition than before. This smoothness on the edges is due to the use of the simple mean contact model. This model is extremely noisy and displays an almost random behavior on edge points.

Despite all this, the main difference between the scans made extensively and the the scans using the Bayesian optimization method is the number of trials involved. On the extensive case the number of samples taken is constant for all objects while for the number of samples taken on the Bayesian scans vary with the object. This variations are explained by the complexity of the object and the non-fixed evolution of the Bayesian method. The number of samples taken with each method for each object is shown in table 4.1. Notice that, for the Bayesian optimization method, several additional

values are displayed. These values represent the number of iterations of the method that were needed for it to find a value that differs from the global maximum by less than 5%, 10% and 20% respectively from left to right.

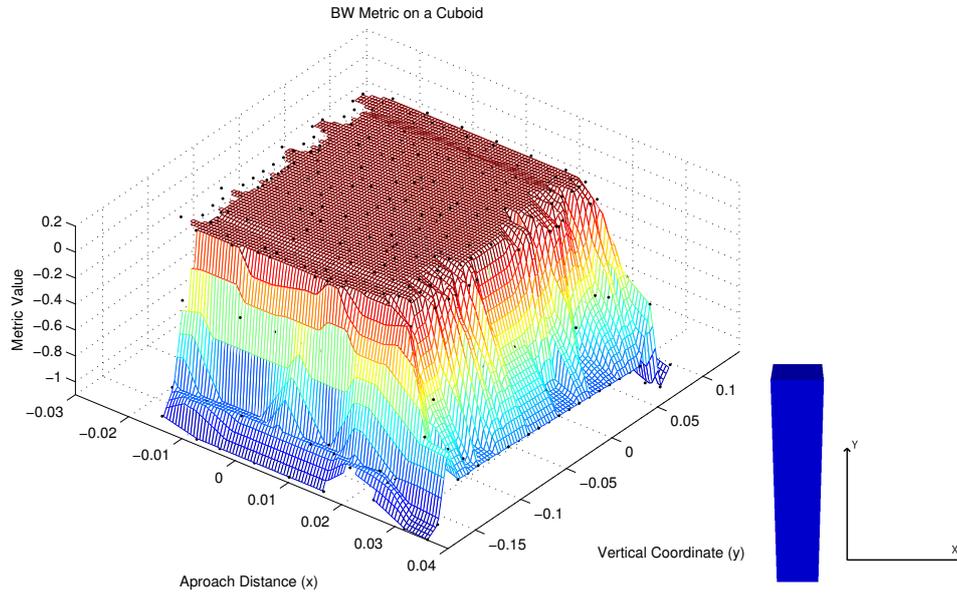


Figure 4.22: *BW* Metric values of grasps in a cuboid, computed using the Bayesian Optimization method.

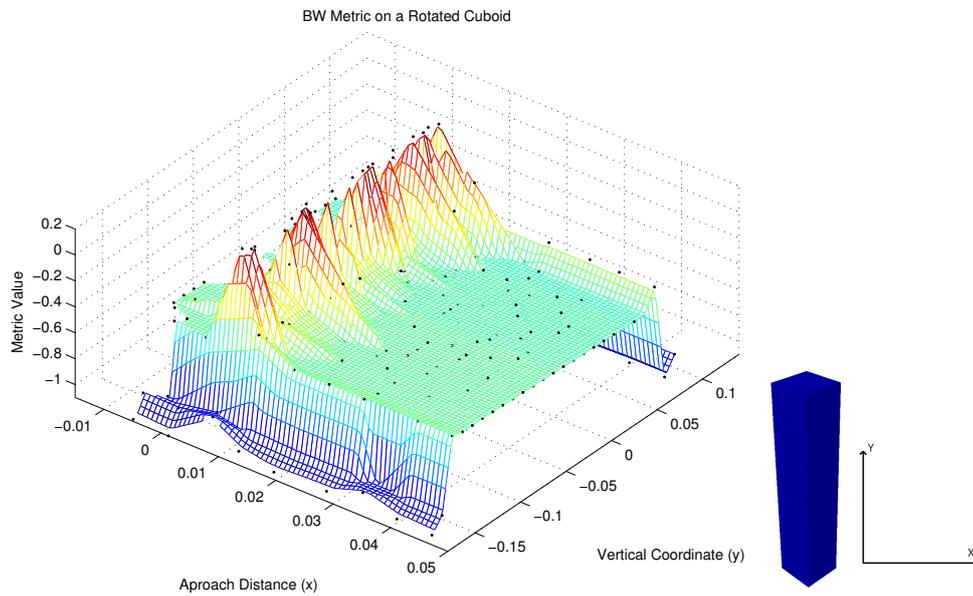


Figure 4.23: *BW* Metric values of grasps in a rotated cuboid, computed using the Bayesian Optimization method.

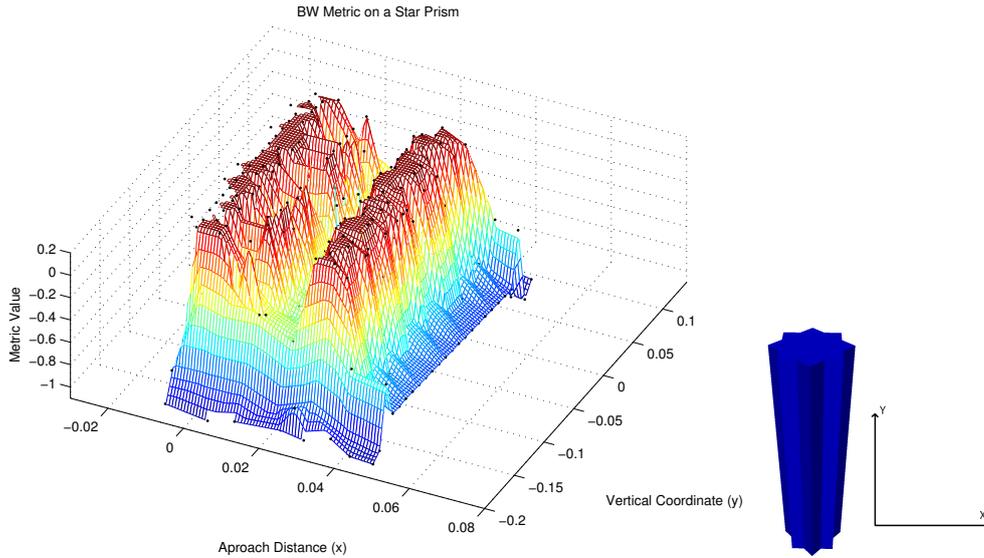


Figure 4.24: *BW* Metric values of grasps in a star prism, computed using the Bayesian Optimization method.

Table 4.1: Number of Samples taken by each method for each object.

Object	Extensive Search	Bayesian Optimisation			
		Total	5%	10%	20%
-	2211	551	14	14	14
Sphere		306	31	19	19
Wine Glass		485	39	20	20
Cylinder		331	55	55	1
Mug		564	1	1	1
Cuboid		388	59	59	59
Rotated Cuboid		402	105	105	31
Star Prism					

Observing table 4.1 one can see that although only stopping when the expected improvement no longer displays interesting values, the Bayesian optimization method could have settled for much less samples and still provide a fairly good approximation of the global maximum.

4.8 Bayesian optimization versus random sampling

In this section the results of a base comparison between the performance of the Bayesian optimization method and random sampling on parameter space are shown. This comparison is done by measuring the evolution of the best value found by each algorithm along a sequence with 100 iterations when performing 2D sampling. For more robust results, we consider the mean evolution along 5 trials.

Fig. 4.25 represents the evolution of the best value found by each method when sampling as sphere. The Bayesian method (black) clearly converges faster than the random sampling method (red) to the maximum function value.

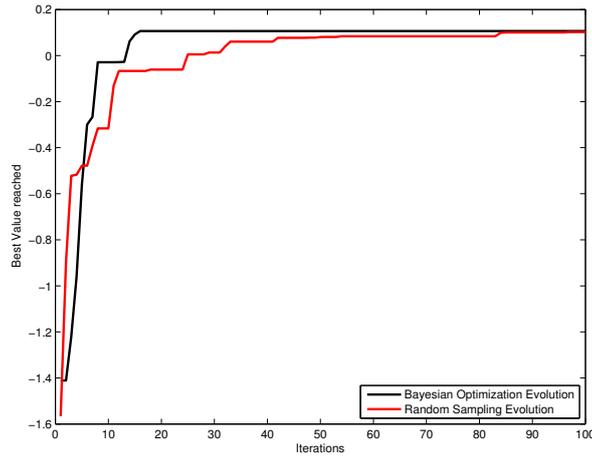


Figure 4.25: Comparing the evolution of best value found when sampling a sphere with the Bayesian approach versus the random sampling method.

To confirm the results obtained by the previous experiment, similar experiments were made using a mug and a star prism instead of the sphere. Fig. 4.26 shows that although the Bayesian method still converges faster to the maximum function value, the random sampling method converges faster to high values of the function.

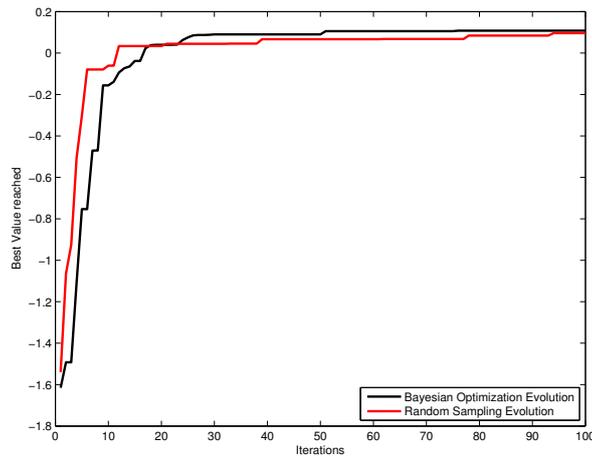


Figure 4.26: Comparing the evolution of best value found when sampling a mug with the Bayesian approach versus the random sampling method.

In Fig. 4.27 the results are similar for both methods.

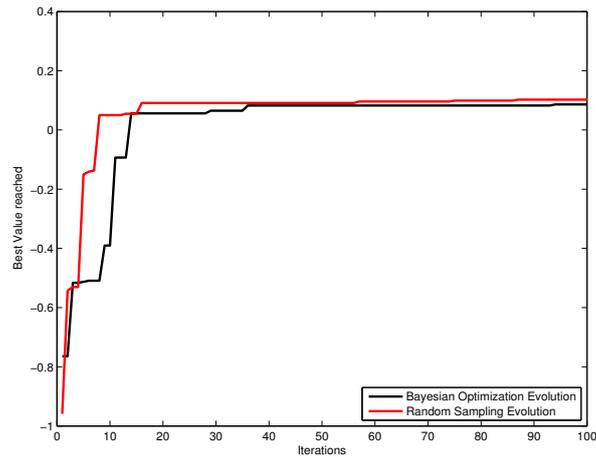


Figure 4.27: Comparing the evolution of best value found when sampling a star prism with the Bayesian approach versus the random sampling method.

We conclude that the Bayesian method converges faster to the maximum value of $f(x)$ when compared with the random sampling method but that for simple and highly symmetrical objects the difference between the two methods is not significant when working in 2-dimensional spaces.

4.9 Results on a real platform

In this section we show the system working in a real platform. Figure 4.29 shows the results of a preliminary integration with the main experimental platform for the HANDLE project, fig. 4.28.



Figure 4.28: Shadow arm and hand that serves as main platform for the HANDLE project (left), image captured by the Kinect mounted on the arm during the course of the experiment.

The experiment consisted of sampling a can along the radial axis. The positive spikes shown in fig. 4.29 represent evaluations of the metric when the grasp is complete. The rest of the data variations reflect the transitions from non existing or partial contact between the fingers and the object in order to

arrive at the final grasp configuration. We emphasize the fact that this experiment is only a preliminary integration where no bayesian methods were used. Despite being preliminary, results are promising in the sense that the metric variation correctly depicts what would be expected for this object shape.

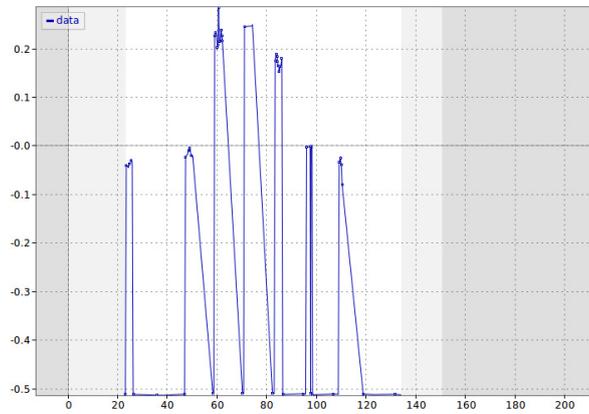


Figure 4.29: Metric values taken during the experiment.

5

Conclusions and Future Work

This work proposed a methodology to compute optimal or close-to-optimal grasps on a diversity of scenarios. We considered both known and unknown objects, simulated or real robotic setups and offline and online usage. Two key issues contributed to the success of the approach. First a new metric to evaluate the quality of grasps was developed to assess both stable and unstable grasps, thus enlarging the domain of application. Whereas classical metrics concentrate on force-closure grasps, we believe that the quantification of “distance” to force-closure can bring important benefits to real-world robot grasping. We have shown in simulation with several objects that our metric can cope with a larger diversity of situations when compared with classical metrics. The fact that its optimal values are invariant to reference frame location opened the possibility for online operation. We have shown that it is possible to create a robotic grasping system that “feels” the object and gradually searches for the optimal grasp, through the use of Bayesian Optimization methods. We showed that these methods are more efficient than extensive or random search of the parameter space. We also showed that they provide a measure of the system’s global uncertainty. This measure can be interpreted as how much do we actually understand the complete system. The only requirement for the use of the system that was described is that the object does not move during the learning process making it extremely attractive for several applications.

Future work will keep this path of research. Extending the parameter space to include more information on the initial position, such as full translation and rotational properties, introducing the hand joint space or simplified representations of it or even introducing measurements on object properties may enrich the method in ways yet to be studied. In real robotic grasping it is often the case that planned grasps cannot be executed perfectly due to several sources of uncertainty and system mis-calibrations. Therefore the execution of an adequate grasp may not be possible at the first contact. A sensible measurement of grasp quality at the first contact or even the matching of this measurement to knowledge attained through simulation may be critical to define a plan of action allowing to achieve a good grasp in a few iterations.

Bibliography

- [1] A. Miller and P. Allen, "Graspit! a versatile simulator for robotic grasping," *Robotics & Automation Magazine, IEEE*, December 2004.
- [2] R. Diankov, "Openrave: A planning architecture for autonomous robotics," *Robotics Institute, Pittsburgh, PA, Tech. Rep.*, July 2008.
- [3] E. Brochu and V. Cora, "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.
- [4] H. Liu, X. Song, J. Bimbo, K. Althoefer, and L. Seneviratne, "Intelligent Fingertip Sensing for Contact Information Identification," in *ASME/IFToMM Int. Conf. Reconfigurable Mechanisms and Robots*, 2012.
- [5] J. Felip and A. Morales, "Robust sensor-based grasp primitive for a three-finger robot hand," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. IROS 2009*. IEEE, 2009, pp. 1811–1816.
- [6] Z. L. R. Murray and S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. Boca Raton, FL: CRC Press, 1994.
- [7] Z. Li and S. Sastry, "Task-oriented optimal grasping by multifingered robot hands," *IEEE Journal of Robotics and Automation*, vol. 4, no. 1, pp. 32–44, 1988.
- [8] C. Ferrari and J. Canny, "Planning optimal grasps," *International Conference on Robotics and Automation*, 1992.
- [9] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," *Robotics & Automation Magazine, IEEE*, pp. 348–353, 2000.
- [10] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [11] D. E. Finkel, "DIRECT Optimization Algorithm User Guide," 2003.

