

# Bayesian Optimization for Robust Robotic Grasping Using a Sensorized Compliant Hand

Juan G. Lechuz-Sierra , Ana Elvira H. Martin , Ashok M. Sundaram , Ruben Martinez-Cantin ,  
and Máximo A. Roa , *Senior Member, IEEE*

**Abstract**—One of the first tasks we learn as children is to grasp objects based on our tactile perception. Incorporating such skill in robots will enable multiple applications, such as increasing flexibility in industrial processes or providing assistance to people with physical disabilities. However, the difficulty lies in adapting the grasping strategies to a large variety of tasks and objects, which can often be unknown. The brute-force solution is to learn new grasps by trial and error, which is inefficient and ineffective. In contrast, Bayesian optimization applies active learning by adding information to the approximation of an optimal grasp. This paper proposes the use of Bayesian optimization techniques to safely perform robotic grasping. We analyze different grasp metrics to provide realistic grasp optimization in a real system including tactile sensors. An experimental evaluation in the robotic system shows the usefulness of the method for performing unknown object grasping even in the presence of noise and uncertainty inherent to a real-world environment.

**Index Terms**—Grasping, force and tactile sensing, learning from experience.

## I. INTRODUCTION

**E**ARLY work on robotic grasp prediction, including the work carried out by Miller et al. in [1], was based on deterministic models with programmed grasping sequences constrained by the need of having a precise geometric and physical model of the object to be grasped. Later works have shifted the focus to more data-driven and adaptive grasping strategies, which are more flexible, robust and reliable, such as the approach by Yang et al. [2] or the Dex-Net system developed in [3] by Mahler et al. However, these approaches are not without drawbacks since they are computationally expensive and require large amounts of data to be trained. On the other

Received 27 May 2024; accepted 19 September 2024. Date of publication 7 October 2024; date of current version 14 October 2024. This article was recommended for publication by Associate Editor S. Sakaino and Editor T. Ogata upon evaluation of the reviewers' comments. This work was supported in part by the EU Horizon Europe Research and Innovation Program Intelliman under Grant 101070136, in part by Spanish Government CaRo under Grant PID2021-125209OB-I00 and Grant PID2020-114819GB-I00, and in part by Aragon Government under Grant DGA-T45\_23R. (*Corresponding author: Juan G. Lechuz-Sierra.*)

Juan G. Lechuz-Sierra and Ruben Martínez-Cantin are with the Instituto de Investigación en Ingeniería de Aragón (I3A) and DIIS, Universidad de Zaragoza, 50009 Zaragoza, Spain (e-mail: jgarcialechuz@unizar.es).

Ana Elvira H. Martin, Ashok M. Sundaram, and Máximo A. Roa are with the German Aerospace Center (DLR), Institute of Robotics and Mechatronics (RM), 82234 Weßling, Germany.

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2024.3475914>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2024.3475914




Fig. 1. Different setups used during experimentation with Bayesian optimization-based grasps: (a) Simulation scenario, optimizing the grasp by moving CLASH hand freely around the object; (b) Simulation of the complete robot model, with motion and collision constraints due to the arm and workbench; (c) Real experimentation environment, including CLASH, a compliant under-actuated hand with tactile sensors on the fingertips.

hand, to enhance the efficiency and adaptability of a grasping method it is crucial to make it transferable to any object, even without previous knowledge of its characteristics. A simple solution to address the grasping of an unknown object could be to learn new grasps by trial-and-error methods, which can be highly inefficient and computationally demanding. In this line, different model-free approaches like the work carried out by Mahler et al. [4] and Yi et al. [5] have been proposed, planning the grasp based on visual and tactile feedback and statistical foundations with greater interpretability than deep learning approaches. Among them, Bayesian Optimization (BO) [6] has the capacity to leverage online grasp memory, making it a powerful tool for creating versatile robotic grasping systems. It is the most sample-efficient global optimization method, reducing the number of trials required to find an optimal grasp, balancing between exploring new configurations and exploiting the learned probabilistic features that describe the process, while handling the uncertainty inherent in robotic perception and control. In particular, working with an approach based solely on tactile detection, with no external sensor for shape recognition, allows us to confront BO with the uncertainty introduced by realistic sensors, for analytically measuring the stability of the grasp. Nevertheless, its sample efficiency makes it suitable to deal with general black-box functions that may include visual feedback to ensure robustness.

In this work, a BO approach based on tactile feedback is used to handle complexity and uncertainty introduced by a real-world setup (Fig. 1(c)), proving the generality of the method to obtain robust and safe grasps when dealing with diverse objects and

end-effectors just by using the knowledge from previous trials. The specific contributions of this paper are:

- An exploration system for robotic grasping in 3D space based on BO, applied to a robotic hand endowed with tactile sensors.
- The development of new heuristics that improve the convergence of the grasp evaluation function, increasing the applicability and task success in real environments.
- Verification of the portability of the method to a real system, evaluating grasp using contact points-based analytic metrics and real tactile sensor measurements.

## II. RELATED WORK

**Robotic grasping** has greatly evolved from only relying on pre-programmed sequences, such as the work by Miller et al. in [1], where the modeling of the object as a shape primitive allows the selection of a starting grasp position. To compensate for the lack of precision and adaptability of such approaches, learning-based methods have lately received increased attention. Data-driven strategies have shown improved robustness through the use of extensive real-world data. This is the case of the Dex-Net system developed by Mahler et al. in [3], where sets of point clouds were used to learn the properties of the object. Although that work focused on reducing the required training data, deep learning methods are still computationally expensive. To reduce this training time, Yan et al. [7] proposed to use explicit ensemble methods. However, deep-learning-based methods fail to generalize to new objects different from those trained on. Model-free approaches have been explored to facilitate the grasping of previously unknown objects. They rely on exploring grasp possibilities and object shape, according to probabilistic foundations. In the case of Mahler et al. [4], grasp planning is addressed in 2D using visual data only. On the other hand, Yi et al. [5] used tactile feedback to reduce the object surface uncertainty.

To improve the efficiency of exploration and grasp sampling, **Bayesian Optimization** has emerged as a promising technique with a sequential strategy for addressing systematic uncertainties while providing an effective solution in many robotics and reinforcement learning problems, such as robotic walking [8], adaptive control [9], and path planning [10]. Its effectiveness in robotic grasping was first studied in the work of Daniel et al. [11], where a human acted as an expert to evaluate grasps. Later, Nogueira et al. [12] introduced the unscented transformation as a system to propagate the noise distribution through decisions, enabling the identification of optimal and safe grasps, with no human needed in the loop. The work was later continued by Castanheira et al. in [13], where a haptic exploration strategy combining BO with a collision penalty heuristic is proposed, accelerating the method's convergence. Another important contribution was made by De Farias et al. in [14], introducing object shape representation and tactile exploration into the grasp planning process, achieving stable grasps with a high probability of success. Recently, our contribution with Herrera et al. [15], compared different parameterizations of the grasp to be used

with BO, as well as BO algorithm variants to obtain multiple grasp solutions. Compared to previous work, this work confronts the probabilistic nature of BO directly with the uncertainty introduced by real tactile sensors, which makes contact points-based metrics less reliable.

**Tactile sensing** has been utilized in intelligent robotics for slip detection [16], manipulation [17] and object properties perception [18], emulating human capabilities by multi-modal signal measurement, such as in the sensors used in this work with tri-axis force measurements. In addition, a tridactyl under-actuated soft hand has been used, which, while increasing the grasping capabilities [19], decreases the measurement accuracy, due to the positional error introduced when computing the analytical metrics. This increases complexity compared to grasping with pincer-like grippers used in other grasping approaches so far.

## III. BAYESIAN OPTIMIZATION

This section explains the mathematical background allowing BO to efficiently tackle the global optimization of black-box functions, e.g. grasping an object. More details about the notation, implementation and pseudo-code can be found in the literature [6], [20].

Formally, BO finds the optimum of an unknown real-valued function  $f : \mathcal{X} \rightarrow \mathbb{R}$ , where  $\mathcal{X} \subset \mathbb{R}^d$  is a compact space, in our case corresponding to the space of parameters describing the end-effector pose. The algorithm selects the best query point  $x_i \in \mathcal{X}$  to evaluate each iteration with outcome  $y_i = f(x_i) + \eta$ , where  $\eta$  is zero-mean noise with variance  $\sigma_\eta^2$ , finding an optimum  $x^*$  so that  $|y^* - y_n|$  is minimum for the considered budget.

BO is composed of two main ingredients. The first one is a **probabilistic surrogate model**, which is a distribution over the family of functions  $P(f)$ , where  $f$  belongs, built by incrementally sampling over  $\mathcal{X}$ . As implemented in BayesOpt [20] we consider the probabilistic model as a Gaussian process  $GP(x|\mu, \sigma^2, \theta)$  with inputs  $x \in \mathcal{X}$ , scalar outputs  $y \in \mathbb{R}$  and an associated kernel or covariance function  $k$  with hyperparameters  $\theta$ , estimated using Monte Carlo Markov Chain (MCMC) algorithm resulting in samples  $\{\theta_i\}_{i=1}^m$ . Given a set of observations  $\mathcal{D}_n = (x_{1:n}, y_{1:n})$  at step  $n$ , then the prediction of the GP  $y_{n+1} = \hat{y}(x_{n+1})$  at a new query point  $x_{n+1}$ , with kernel  $k_i$  conditioned on the  $i$ -th hyperparameter sample  $k_i = k(\cdot, \cdot | \theta_i)$  is normally distributed,  $\hat{y}(x_{n+1}) \sim \sum_{i=1}^m \mathcal{N}(\mu_i, \sigma_i^2 | x_{n+1})$ , where:

$$\mu_i(x_{n+1}) = k_i(x_{n+1}, X)K_i^{-1}y \quad (1)$$

$$\sigma_i^2(x_{n+1}) = k_i(x_{n+1}, x_{n+1}) - k_i(x_{n+1}, X)^T K_i^{-1} k_i(X, x_{n+1}) \quad (2)$$

where  $k_i(x_{n+1}, X) = [k_i(x_{n+1}, x_j)]_{x_j \in X}$  and  $K = [k_i(x_j, x_l)]_{x_j, x_l \in X} + \mathbf{I}\sigma_n^2$  being  $\sigma_n^2$  the noise, with  $\mu_i$  and  $\sigma_i^2$  representing the prediction and uncertainty of the model at  $x_{n+1}$ .

The second one is an **acquisition function**  $\alpha(x, p(f))$ . It considers the predictive distribution for each point in  $\mathcal{X}$  to select the next point at each iteration. The criteria used to perform this task is the Expected Improvement criterion (EI), which, given a set of observations  $\mathcal{D}_n = (X, y)$  at step  $n$  and considering the

expectation over the predictive distribution, can be computed as:

$$\begin{aligned} EI(x) &= \mathbb{E}_{(\hat{y})|D_n, \theta, x}[\max(0, \hat{y}(x) - \rho)] \\ &= \sum_{i=1}^m [(\mu_i(x) - \rho)\Phi(z_i) + \sigma_i(x)\phi(z_i)] \end{aligned} \quad (3)$$

where  $\rho$  is an incumbent value, i.e.  $\rho = \max(y_{1:n})$ ,  $\phi$  and  $\Phi$  correspond to the probability and cumulative density functions respectively, and  $z_i = (\rho - \mu_i(x))/\sigma_i(x)$ . The pair  $(\mu_i, \sigma_i^2)$  is the mean and variance of the predictive distribution. EI can be maximized by optimizing two main elements of (3), which are  $(\mu_i(x) - \rho)$ , for exploitation, and  $\sigma_i(x)$ , for exploration.

#### IV. GRASP EVALUATION

Bringing the iterative process of BO to robotic grasping terms, in the first step of the optimization, the grasping process is fed with an initial set of samples  $\mathcal{X}$  of the grasp pose. For the selected set of samples, the grasping process takes place, updating the probabilistic model and the acquisition function with the grasp quality computed for each of the samples. In the remaining iterations, the pose of the hand that maximizes the acquisition function is computed and each time a new target pose is received, the hand is placed in the new location using position control, and force-controlled to grasp the object, and a new evaluation of the function developed below is carried out. We optimize  $\{x, y, z\}$  coordinates of the hand palm based on the bounding box of the object plus the finger size in each direction. We also optimize the roll angle. Pitch and yaw are computed to align with the object center [21]. It is important to mention that this 4D problem is higher than in previous research with 2D [12] or 3D [13] parametrizations.

The following section introduces all the processes that allow the method to be transferred from a limited simulation (Fig. 1(a)) to the real system (setups in Fig. 1(b) and (c)).

##### A. Reachability and Collisions

For finding potential grasp candidates, we freely sample the space around the object to retrieve potential grasping poses. However, this means that some hand-palm poses and robot configurations can be unfeasible for different reasons.

As a first step to evaluate the grasping poses provided by the acquisition function of the BO, the reachability of such pose is checked by comparing it with a Capability Map [22], whose cells have an associated index representing the dexterity of the robot in that location. Once this check has been passed, we compute the robot inverse kinematics to provide a specific robot joint configuration.

Another type of unfeasible grasping pose occurs when the robot is in collision with either the object or the workbench. A possible collision with the workbench is handled by classifying the grasp as invalid, assigning a grasp quality equal to 0, thus moving the acquisition function away from poses similar to this one. For collisions between hand and object, an Approximation Reward ( $AR$ ) is used to differentiate them from collisions with the workbench or from poses that do not even touch the object, thus guiding the search toward regions that are worth exploring.

Inspired by the Collision Penalty ( $CP$ ) in [13], this  $AR$  brings the optimization closer to the object while avoiding those poses in which object and hand collide, assigning a higher value to the lower level of penetration of the hand into the object, as follows:

$$AR(n_j) = \begin{cases} e^{-\lambda n_j}, & \text{if } n_j > 0 \\ 0, & \text{if } n_j = 0 \end{cases} \quad (4)$$

where  $\lambda$  acts as a decay rate constant and  $n_j$  is the number of hand links colliding with the object in the pose sampled by BO. When there are no collisions between hand and object and  $n_j = 0$ ,  $AR$  is also 0. This kind of heuristic does not need precise knowledge about object shape, and the pre-calculated collisions can be obtained using the robot model and a bounding box, which can be estimated in a real environment from a single point of view.

##### B. Absence of Force Closure

In a typical grasp planning approach, when the grasp generated after the hand closes around the object does not lead to force closure, the grasp quality is equal to 0 to feed the optimization. However, taking as inspiration the  $AR$ , a Contact Reward ( $CR$ ) has been implemented to improve the method's convergence. This  $CR$  is defined as:

$$CR(n_c) = 1 - e^{-\lambda n_c} \quad (5)$$

where  $n_c$  is the number of contact points at the hand's fingertips and  $\lambda$  has the same purpose as in (4). In this way, the acquisition function is directed towards positions where there is a higher number of fingertip contacts with the object.

##### C. Grasp Quality Computation

When an iteration results in a grasp with force closure, the quality of the grasp is computed. Several works [23], [24] have concluded that the use of a combined metric is advantageous for finding an optimal realistic grasp. Based on these works, the following metrics were studied: to consider the hand configuration, the Uniformity of the transformation in the velocity domain from the finger joints to the object is computed using the condition number of the Jacobian matrix ( $Q_{uni}$ ) [25]; the Grasp Isotropy Index ( $Q_{iso}$ ) [26] is computed from the singular values of the Grasp matrix to indicate an isotropic contribution of contact forces to the applied wrench; further analysis of the contact forces is performed using the wrenches that the grasp can resist, known as Grasp Wrench Space (GWS), by considering the Largest-minimum resisted wrench ( $Q_\epsilon$ ) [27] and the Volume of this GWS ( $Q_v$ ) [28]. From experiments in Section V-A and V-C, weights were assigned to each metric, obtaining the combined metric in (6), where  $\sum_{i=1}^4 w_i = 1$ . Weights can also be optimized by other methods making  $Q_m$  more robust.

$$Q_m = w_1 * Q_{iso} + w_2 * Q_\epsilon + w_3 * Q_v + w_4 * Q_{uni} \quad (6)$$

$Q_m$  also aims to compensate metrics with worse convergence with those with better convergence and to compensate those more influenced by the uncertainty of the real environment. The




Fig. 2. New metrics proposed: (a) Approximation Reward in two grasps with collision, from higher  $AR$  (left) to lower  $AR$  (right). (b) Contact Reward in two grasps with finger contact points, from higher  $CR$  (left) to lower  $CR$  (right).



Fig. 3. Objects from the YCB set used for experimentation, both in simulation and in the real environment. They show variety in shapes, weights, sizes, symmetry or the position of the center of mass.

TABLE I  
ABLATION STUDY. MEAN AND MAXIMUM  $y_{simple}$  OUTCOMES OBTAINED GRASPING A MUSTARD BOTTLE WITH DIFFERENT GRASP EVALUATION FUNCTIONS

		$y_{simple}$		Cas.[13] FM	$y_{AR}$ FM	$y_{gr}$ FM
		FH	FM			
Mustard bottle	Mean	0.599	0.318	0.546	0.562	0.657
	Max.	0.967	0.518	0.815	0.901	0.935

FH: free hand; FM: full model.

resulting evaluation function is of the form:

$$y_{gr} = Q_c * (Q_m + \alpha(AR + CR)) \quad (7)$$

where  $Q_c$  is a binary index indicating whether the pose is feasible or not. The  $\alpha$  coefficient has been added to regulate the range of values spanned by  $AR$  and  $CR$ , depending on the magnitude of the values obtained when calculating the grasp quality metrics. In this way, a higher outcome is always given for a pose in which the object is successfully grasped. Note that both the  $AR$  and  $CR$  are intended to guide the acquisition function without hindering the approximation to a robust grasp.

## V. EXPERIMENTAL RESULTS

In this section, we present the experiments conducted in simulation and with the real robot. The simulation was implemented using PyBullet, and the objects considered (Fig. 3) were selected from the publicly available YCB object set [29].

### A. Ablation Study

For establishing a baseline for comparison, we use the free-hand simulation from [15], tested in the CLASH hand with  $w_2 = 1$  for computing  $Q_m$  with (6). The baseline is presented with a dashed line in Fig. 4 and in Table I, showing BO's ease in finding feasible grasps with great outcome values. The corresponding evaluation function has the form presented in (8), where  $Q_f$  is a binary index indicating the force closure condition in this




Fig. 4. Comparison of the evolution of the  $y_{simple}$  outcome, obtained using different implementations of the grasp evaluation function.

experimental case.

$$y_{simple} = Q_c * Q_f * Q_m(w_2 = 1) \quad (8)$$

The following experiments show a comparison between the results obtained using different implementations of the grasp evaluation function on the complete robot model simulation, shown in Fig. 1(b). As in the previous case, we use  $w_2 = 1$ , with 20 initial samples and 50 BO iterations to optimize each evaluation function. The plots in Fig. 4 include a comparison of the outcomes obtained in (8), as a common metric, with the grasps found by each evaluation function in 10 optimizations, and the shaded area represents the 95% confidence interval.

When adding the arm and the workbench in the simulation, the obtained metric values are much lower than in the case of the free-hand simulation, as observed in Table I and the orange line in Fig. 4. The reachability of the grasping pose and the collisions detected when including the arm and the workbench result in a higher number of grasping poses evaluated as 0. We also added the  $AR$  value to the evaluation function, leading to (9). Throughout all the experimentation, the tuning parameter was set to  $\lambda = 0.1$ , and a coefficient  $\alpha = 0.1$  was used, resulting in  $AR$  values between 0 and 0.1, being substantially lower than those provided by the grasp metric, which has been normalized between 0 and 1.

$$y_{AR} = Q_c * Q_f * (Q_m(w_2 = 1) + \alpha AR) \quad (9)$$

Table I shows how both mean and maximum outcomes increase with the addition of  $AR$ , and in Fig. 4 we can see how active learning reaches much higher values already in the first iterations. The comparison between the results obtained by applying the  $CP$  [13] and the  $AR$  proposed in this work demonstrates how assigning a higher value to  $AR$  the lower the degree of collision with the object drives the acquisition function towards near-object poses, improving convergence while preventing collisions.

In the next step of the experimentation, we added  $CR$ , resulting in (7), again with  $w_2 = 1$ . The results of Fig. 4 and Table I show the greater convergence of the method when including this coefficient. Table II shows the results obtained on more complex objects, for grasps that are evaluated with a  $y_{gr} > 0.5$ .

To continue with the full model simulation experiments, the performance of the method using each metric was tested. Table III shows the time for each metric to achieve a  $y_{gr}$  outcome

TABLE II  
FULL MODEL SIMULATION. MEAN AND MAXIMUM OUTCOMES OBTAINED BY GRASPING DIFFERENT OBJECTS USING  $w_2 = \ln y_{gr}$

	Mustard bottle	Mug	Power drill
Mean	0.6572	0.6062	0.4964
Max.	0.9345	0.8713	0.6246

TABLE III  
FULL MODEL SIMULATION: CONVERGENCE OF EACH METRIC UNTIL  $y_{gr} > 0.5$

		$w_1 = \mathbf{1}$	$w_2 = \mathbf{1}$	$w_3 = \mathbf{1}$	$w_4 = \mathbf{1}$
Mustard bottle	time (s)	12.52	31.00	55.46	18.77
	$y_{gr} < 0.5$	0/10	3/10	5/10	0/10
Mug	time (s)	21.31	32.93	46.13	14.94
	$y_{gr} < 0.5$	0/10	4/10	5/10	0/10
Power drill	time (s)	27.16	38.95	38.15	12.02
	$y_{gr} < 0.5$	3/10	6/10	6/10	0/10




Fig. 5. Some of the best grasps obtained with each metric in simulation, grasping the mustard bottle, the mug, and the power drill. From left to right in the cases of the mustard bottle and the mug:  $Q_{iso}$  ( $w_1 = 1$ ),  $Q_\epsilon$  ( $w_2 = 1$ ),  $Q_v$  ( $w_3 = 1$ ) and  $Q_{uni}$  ( $w_4 = 1$ ). In the case of the power drill:  $Q_{iso}$  and  $Q_\epsilon$  in the upper part,  $Q_v$  and  $Q_{uni}$  at the bottom.

greater than 0.5, and the number of times that this minimum outcome was not reached. Considering metrics normalization, we choose this threshold based on the literature [24], providing a logical division between above and below-average performances and proving to be correlated with a reasonable success rate. It could be increased to provide more demanding applications. However, complementing this approach with experimental validation is key for a meaningful evaluation. We evaluated the grasps with the function provided in (7), assigning a weight of 1 to the corresponding grasp metric ( $Q_{iso}$ ,  $Q_\epsilon$ ,  $Q_v$ ,  $Q_{uni}$ ). The same number of iterations as in the previous experiments was used.

The average convergence time obtained with  $Q_{iso}$  and  $Q_{uni}$  is much lower than the time obtained when using  $Q_\epsilon$  and  $Q_v$ . Note that these last two metrics did not reach the minimum  $y_{gr}$  outcome a greater number of times, while  $Q_{uni}$  is the only metric that reached this minimum value in the totality of the experiments. On the other hand,  $Q_v$  was the metric with the worst performance in obtaining an optimal grasp on each object.  $Q_{iso}$  and  $Q_{uni}$  are ratios of the minimum and maximum singular values of their respective matrices and are therefore unitless. When normalized, their range of values will be lower and closer to 1 if the differences between the maximum and minimum singular values on each case remain similar. On the other hand,  $Q_\epsilon$  and  $Q_v$  are absolute measures from the wrench space. Their variation can be more accentuated because of the influence of small disturbances in the position or direction of the forces exerted on the object. Fig. 5 shows some of the best

grasps obtained by the optimization of (7) during the described experiments, grasping each considered object.

### B. Real-World Setup

Our robot setup, presented in Fig. 1(c) consists of a 7 degrees of freedom LWR arm [30] and a three-fingered CLASH hand [31], with variable stiffness in the fingers and including a grid of  $4 \times 4$  haptic sensors XELA uSkin Patch [32] on the fingertips, measuring force in three axes. Together with the finger coupling, CLASH variable stiffness allows greater adaptability in grasps with different types of objects, in which a lower tendon pretension can prevent the slippage of a curved object or deformation of very delicate objects when pressure is exerted. The contact point analysis requires the computation of the pose of each sensor using the kinematics of the arm-hand assembly. For this purpose, the hand has been considered fully actuated, which does not correspond to reality. Carrying out the real experimentation with an under-actuated hand adds more error to the calculations performed to evaluate the grasp. The tendon-driven fingers make it difficult to obtain an accurate joint position, which translates into another source of uncertainty that BO has to handle. The CLASH hand provides interfaces to control the finger's stiffness and joint positions independently, and the LWR arm provides interfaces for both Cartesian and joint impedance control. These controllers run on a Linux real-time machine, while other software components, including BayesOpt and the Pybullet simulation for collision avoidance, are executed in a separate Linux PC with an Intel Xeon E5-1630 v4 CPU @ 3.70GHz and an NVIDIA Quadro K620 GPU.

One restriction of performing movements on the real arm is that most grasping poses result in collisions when the hand approaches the target directly. To solve this type of collision a pre-grasp position is established, several centimeters away from the target position in the local Z-axis of the hand, from which the hand approaches the object safely every iteration. Additional information on the implementation problems addressed can be found in [21].

### C. Simulation vs. Real Robot

To verify the portability of the method to the real robot, in this experiment different grasping poses were obtained in simulation using each of the implemented metrics and three different objects, and then tested on the real robot. The number of successful grasps in 5 optimizations, using the same number of iterations as in the previous experiments, was obtained for each metric and object, to compare the metrics' effectiveness. The outcomes have been plotted in Fig. 6 with their 95% confidence intervals, representing the learning differences of each metric. To evaluate grasp success, the object was lifted several centimeters from the table and held in that position for a period of 5 seconds. A binary score (success or fail) was recorded for each optimal grasp pose. A grasp execution was considered a failure if the object fell or slipped during its lift motion, touching the workbench. The results obtained in this experiment are summarized in Table IV.

We can observe mean  $y_{gr}$  outcomes greater than 0.5 in most of the cases. This is especially interesting in the case of  $Q_v$

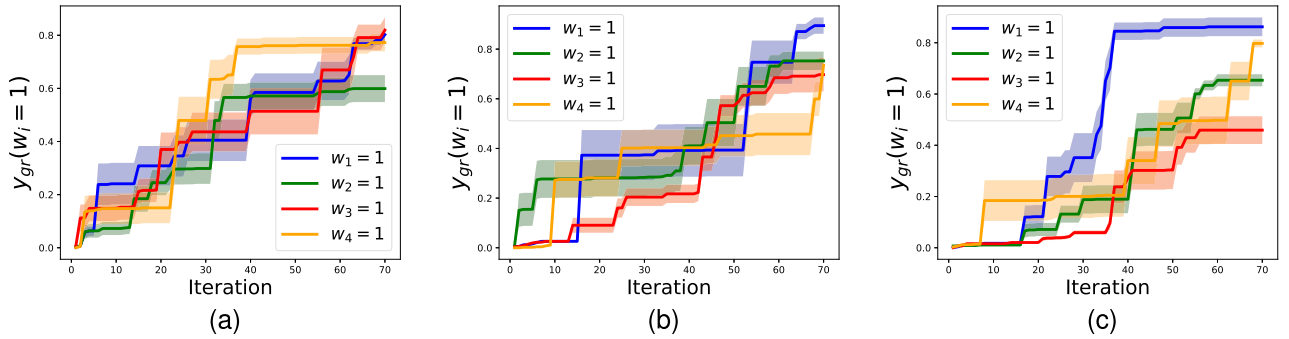


Fig. 6. Comparison of the evolution of the grasp quality, to test optimal results in the real setup. Each line represents the outcome of (7), assigning a value of 1 to the corresponding weight of each metric and grasping different objects: (a) Pringles can. (b) Mug. (c) Lego toy.




Fig. 7. Grasps optimized in simulation with each metric and tested on the real robot. Figures (a)–(d) correspond to successful grasps holding the Pringles can, using  $Q_{iso}$ ,  $Q_{\epsilon}$ ,  $Q_v$ ,  $Q_{uni}$ , in this order. Figures (e)–(h) and (i)–(l) correspond to successful grasps holding the mug and Pringles can respectively using the same metrics order. Figures (m)–(p) correspond to failed grasps due to slippage or mislocation.

TABLE IV  
MEAN OUTCOMES AND SUCCESSFUL GRASPS FOR EACH METRIC, OPTIMIZED IN SIMULATION AND TESTED IN THE REAL ROBOT

		$w_1 = \mathbf{1}$	$w_2 = \mathbf{1}$	$w_3 = \mathbf{1}$	$w_4 = \mathbf{1}$
Pringles	Mean outcome	0.8023	0.6991	0.8913	0.7722
	Success	3/5	3/5	1/5	3/5
Mug	Mean outcome	0.8943	0.7526	0.6972	0.7354
	Success	2/5	4/5	3/5	4/5
Lego	Mean outcome	0.8623	0.6537	0.4590	0.7974
	Success	3/5	2/5	1/5	2/5
Total Success		8/15	9/15	5/15	9/15

( $w_3 = 1$ ), which only obtained 33.3% effective grasps in the total evaluated. Both  $Q_v$  and  $Q_{\epsilon}$  are metrics based on GWS properties, which makes them very sensitive to positional errors [33] becoming less accurate when working with sparse and noisy sensor information from the real robot. However,  $Q_{\epsilon}$  is more effective, since it focuses on the maximum force that the grasp can resist and not on the whole wrench space. As seen in Fig. 6, the more complex shape of the object slows down the convergence. However, in most of the cases, the metrics reach an average  $y_{gr}$  outcome greater than 0.5, and the percentage of effective grasps is 65% of the 20 performed per object. In this figure, the blue line corresponding to  $Q_{iso}$  ( $w_1 = 1$ ) reaches the highest outcomes, however, only 2 of the 5 grasps performed

succeeded in lifting the mug (Fig. 6(b)). This indicates a more symmetrical distribution of forces and torques around the object, as in Fig. 7(e) and (g), is not always a sufficient condition to obtain an optimal grasp when faced with the uncertainty introduced by positional errors and an unknown friction coefficient. Although Fig. 6(a) shows very similar learning in all metrics when grasping the Pringles can, in the optimal grasps shown in Fig. 7 we can observe certain trends related to the computation of each metric. While the other metrics focused on exploring the top of the object,  $Q_{uni}$  ( $w_4 = 1$ ) found side grasps in which the hand grasped the can very robustly (Fig. 7(d)). Since CLASH only has tactile sensors on the fingertips, when the Pringles are grasped from above in a sufficiently symmetrical configuration of the fingers, the direction of the contact forces results in a stable grasp according to  $Q_v$ ,  $Q_{iso}$  and  $Q_{\epsilon}$ . However, in our setup, a grasp with the hand wrapped around the can has similar contact points to a grasp touching the can from the side. Considering contact points only on the distal phalanges, and not in the proximals too, causes information to be lost in the analysis. Since objects are not static in the real world, Fig. 7(m) and (n) show grasps in which the object slipped from the hand when pressure was exerted. Fig. 7(o) and (p) show cases where the hand was located far away from the object and the force closure calculation was wrong. On the other hand, Fig. 7(f) shows a case in which the object's movement when exerting pressure favored the grasp since the mug was pushed inside the hand. Fig. 7(h) is another example of how the result obtained by  $Q_{uni}$  wrapped around the

TABLE V  
OPTIMIZATION AND TESTING IN THE REAL ENVIRONMENT

	Mean outcome	Best outcome	Success over 5	Feasible grasp found
Pringles	0.6436	0.8423	3/5	5/5
Mug	0.7189	0.8579	4/5	5/5

Mean  $y_{gr}$  outcomes and number of successful grasps.




Fig. 8. Grasps optimizing  $y_{grasp}$  directly in the real robot. Figures (a) and (b) show successful grasp in different locations holding the Pringles can. Figures (c) and (d) show failed grasps due to errors in force closure estimation, as well as Figure (h) for the case of the mug. Figures (e)–(g) show successful grasps in the mug, varying the roll angle.

mug to a greater extent. In the case of the Lego toy, two types of successful grasps can be observed: the ones in which one or more of the fingers were stuck on the toy’s head, such as those shown in Fig. 7(j) and (l), and those grasps in which the fingers made contact on the edges of the object as shown in Fig. 7(i) and (k). It can also be seen how variable stiffness has increased adaptability to curved surfaces, as in Fig. 7(c) and (h).

#### D. Optimization in Real Environment

In the following experiments, 5 initial samples and 25 BO iterations were performed directly on the real robot, using as evaluation function (7). The weights in (6) were computed as the successful grasps obtained with each corresponding metric in Section V-C, divided by the total number of successful grasps. Collisions and reachability were computed in simulation but the quality of the grasp was obtained from the sensors and measurements provided by the robotic system. Table V shows the number of successful grasps in 5 optimizations and the outcomes obtained, which are also represented in Fig. 9 with their 95% confidence intervals.

In the first optimizations, the importance of the initial samples to conform a well-informed prior can be observed, especially when a reduced number of iterations is considered. In the case of the Pringles can, when a high outcome side grasp was found in the initial samples, the search was focused on that area, more related to greater outcomes in  $Q_{uni}$ , resulting in a grasp like the one in Fig. 8(a). High outcomes in the upper area of the Pringles focused the rest of the search on grasps related to  $Q_e$  or  $Q_{iso}$ , like the one in Fig. 8(b). Although the search space is unconstrained, addressing the learning with heuristics focused on improving the prior can be a good solution for the convergence of the method in areas of interest.

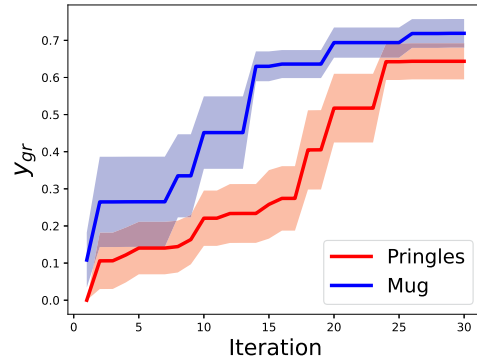


Fig. 9. Comparison of the evolution of the  $y_{gr}$  outcome, optimizing the grasp directly in the real environment.

When working with an underactuated hand of variable stiffness, conventional analytical metrics become more inaccurate as they do not take into account positional errors due to differences in tendon pretension. The force closure computation proved to be insufficient to deal with these errors. Several false good grasps were observed throughout the iterations, two of which are shown in Fig. 8(c) and (d). In general, the use of the GWS to classify the grasp has been proven to have a considerable error due to the uncertainty of the measurements taken for its calculation. Except for  $Q_e$ , none of the metrics considered guarantees the force closure condition in the grasp. Since  $Q_{iso}$  and  $Q_{uni}$  by their definition give rise to more symmetric contact point distributions surrounding the object,  $Q_v$  is the one with the highest risk of obtaining high values for ineffective grasps such as in Fig. 8(h). Even so, Bayesian optimization has been shown to deal positively with these difficulties, achieving a 70% success rate in the optimization of the grasp performed completely on the real robot, as indicated in Table V. Note also that although in 30% of the cases the optimal grasp was not successful, in all cases BayesOpt found other feasible and robust grasps assigned with lower metric values. However, the advantages and disadvantages of each of the metrics have been slightly compensated by using  $Q_m$  to observe a greater variety in the obtained grasps.

Fig. 9 shows faster learning, with higher outcomes for the mug case. The size of the Pringles can, especially in its Z-axis, reduces the convergence of the method but allows obtaining lateral clamp-like grasps. The exploration of the mug gave results that were as successful as those observed in Fig. 8(e) to (g), with more significant variations in the roll angle.

Although the results obtained in [34] give evidence of the effectiveness of their method, they are subject to a prior learning several orders of magnitude higher than the number of trials used in this work to obtain a robust grasp. Moreover, in our case the learning is performed online for any new object. On the other hand, the grasp planner proposed in [3] presents very competitive planning times of less than one second, but which, again, need prior training with large datasets directly affecting the accuracy of the method. The convergence times presented in Table III depend on both software and hardware experimental conditions and still do not exceed one minute of learning in any case, improving also the results obtained with other model-free approaches such as the one presented in [4].

## VI. CONCLUSION

This letter has presented a practical application of BO to the tactile exploration of objects in real environments. It has also shown the effectiveness of BO for grasping diverse and previously unknown objects in a robust and secure manner, using tactile exploration with real-life noise conditions. By conducting comprehensive experiments in simulated and real robotic environments, we have gained valuable insights into the adaptability and performance of this approach in a real setting, accounting for uncertainties and technical complexities.

The use of BO allows the addition of different metrics and heuristics to the grasp evaluation function to deal with diverse environments, without compromising the number of iterations needed. Among these heuristics, an Approximation Reward  $AR$ , a variation to the classical Collision Penalty  $CP$ , has been proposed to force the search algorithm to move away from possible collision, and a Contact Reward  $CR$  has also been included to bring the search closer to more feasible grasping configurations, thus accelerating convergence.

Our consideration of various grasp quality metrics has shed light on their practical utility when applied to a real system, and how they affect the results obtained in BO, which worked significantly well when using noisy measurements to evaluate each grasp trial.

The use of a variable stiffness hand has been beneficial since its great adaptability allows it to grasp curved surfaces robustly while avoiding slipping. Our next steps include the consideration of metrics and a grasp parameterization that better exploits the characteristics of a variable stiffness hand.

## REFERENCES

- [1] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, "Automatic grasp planning using shape primitives," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2003, pp. 1824–1829.
- [2] M. Yang et al., "Sim-to-real model-based and model-free deep reinforcement learning for tactile pushing," *IEEE Robot. Automat. Lett.*, vol. 8, no. 9, pp. 5480–5487, Sep. 2023.
- [3] J. Mahler et al., "Dex-Net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *Robot.: Sci. Syst.*, 2017.
- [4] J. Mahler et al., "GP-GPIS-OPT: Grasp planning with shape uncertainty using Gaussian process implicit surfaces and sequential convex programming," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 4919–4926.
- [5] Z. Yi et al., "Active tactile object exploration with Gaussian processes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 4925–4930.
- [6] E. Brochu, V. Cora, and N. de Freitas, "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, University of British Columbia. Department of Computer Science, 2009.
- [7] G. Yan, A. Schmitz, S. Funabashi, S. Somlor, T. P. Tomo, and S. Sugano, "A robotic grasping state perception framework with multi-phase tactile information and ensemble learning," *IEEE Robot. Automat. Lett.*, vol. 7, no. 3, pp. 6822–6829, Jul. 2022.
- [8] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, "Bayesian optimization for learning gaits under uncertainty: An experimental comparison on a dynamic bipedal walker," *Ann. Math. Artif. Intell.*, vol. 76, pp. 5–23, 2016.
- [9] S. R. Kuindersma, R. A. Grupen, and A. G. Barto, "Variable risk control via stochastic optimization," *Int. J. Robot. Res.*, vol. 32, no. 7, pp. 806–825, 2013.
- [10] R. Martinez-Cantin, N. De Freitas, E. Brochu, J. Castellanos, and A. Doucet, "A Bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot," *Auton. Robots*, vol. 27, pp. 93–103, 2009.
- [11] C. Daniel, O. Kroemer, M. Viering, J. Metz, and J. Peters, "Active reward learning with a novel acquisition function," *Auton. Robots*, vol. 39, pp. 389–405, 2015.
- [12] J. Nogueira, R. Martinez-Cantin, A. Bernardino, and L. Jamone, "Unscented Bayesian optimization for safe robot grasping," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 1967–1972.
- [13] 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 *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 1643–1648.
- [14] C. de Farias, N. Marturi, R. Stolkin, and Y. Bekiroglu, "Simultaneous tactile exploration and grasp refinement for unknown objects," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 3349–3356, Apr. 2021.
- [15] I. Herrera Seara, J. García-Lechuz Sierra, J. Garcia-Barcos, and R. Martinez-Cantin, "Optimización Bayesiana multisolución para la exploración eficiente de agarres robóticos," in *Proc. XLIII Jornadas de Automática*, 2022, pp. 714–720.
- [16] R. A. Romeo and L. Zollo, "Methods and sensors for slip detection in robotics: A survey," *IEEE Access*, vol. 8, pp. 73027–73050, 2020.
- [17] A. Yamaguchi and C. G. Atkeson, "Recent progress in tactile sensing and sensors for robotic manipulation: Can we turn tactile sensing into vision?," *Adv. Robot.*, vol. 33, no. 14, pp. 661–673, 2019.
- [18] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, "Robotic tactile perception of object properties: A review," *Mechatronics*, vol. 48, pp. 54–67, 2017.
- [19] M. Bonilla et al., "Grasping with soft hands," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, 2014, pp. 581–587.
- [20] R. Martinez-Cantin, "BayesOpt: A Bayesian optimization library for non-linear optimization, experimental design and bandits," *J. Mach. Learn. Res.*, vol. 15, pp. 3915–3919, 2014.
- [21] J. Garcia-Lechuz Sierra, "Bayesian optimization for robust robotic grasping," M.S. thesis, Tech. Univ. of Munich (TUM), Sep. 2023. [Online]. Available: <https://elib.dlr.de/197869/>
- [22] O. Porges, T. Stouraitis, C. Borst, and M. A. Roa, "Reachability and capability analysis for manipulation tasks," in *Proc. ROBOT2013: 1st Iberian Robot. Conf.: Adv. Robot.*, 2014, pp. 703–718.
- [23] M. A. Roa and R. Suárez, "Grasp quality measures: Review and performance," *Auton. Robots*, vol. 38, pp. 65–88, 2015.
- [24] B. León, C. Rubert, J. Sancho-Bru, and A. Morales, "Characterization of grasp quality measures for evaluating robotic hands prehension," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2014, pp. 3688–3693.
- [25] J. K. Salisbury and J. J. Craig, "Articulated hands: Force control and kinematic issues," *Int. J. Robot. Res.*, vol. 1, pp. 4–17, 1982.
- [26] B.-H. Kim, S.-R. Oh, B.-J. Yi, and I. H. Suh, "Optimal grasping based on non-dimensionalized performance indices," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2001, pp. 949–956.
- [27] C. Ferrari and J. F. Canny, "Planning optimal grasps," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1992, pp. 2290–2295.
- [28] Z. Li and S. S. Sastry, "Task-oriented optimal grasping by multifingered robot hands," *IEEE J. Robot. Automat.*, vol. 4, no. 1, pp. 32–44, Feb. 1988.
- [29] 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 Robot. Automat. Mag.*, vol. 22, no. 3, pp. 36–52, Sep. 2015.
- [30] G. Hirzinger and A. Albu-Schaeffer, "Light-weight robots," *Scholarpedia*, vol. 3, no. 4, 2008, Art. no. 3889.
- [31] W. Friedl and M. A. Roa, "CLASH—A compliant sensorized hand for handling delicate objects," *Front. Robot. AI*, vol. 6, 2020, Art. no. 138.
- [32] XELA Robotics, "uSPa 44 patch sensor series," Accessed: Jun. 26, 2023. [Online]. Available: <https://www.xelarobotics.com/sensor-collection/uspa-44>
- [33] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2012, pp. 557–562.
- [34] S. Joshi, S. Kumra, and F. Sahin, "Robotic grasping using deep reinforcement learning," in *Proc. IEEE 16th Int. Conf. Automat. Sci. Eng.*, 2020, pp. 1461–1466.