



OPEN ACCESS

EDITED BY

Mohammadreza Davoodi,
University of Memphis, United States

REVIEWED BY

Loris Roveda,
Dalle Molle Institute for Artificial Intelligence
Research, Switzerland
Yazdan Batmani,
University of Kurdistan, Iran
Hamid Hafezi,
Herff College of Engineering, University of
Memphis, Memphis, United States in
collaboration with reviewer YB

*CORRESPONDENCE

Frederik Vinter-Hviid,
✉ fv@ropca.com

RECEIVED 25 November 2023

ACCEPTED 07 March 2024

PUBLISHED 29 April 2024

CITATION

Vinter-Hviid F, Sloth C, Savarimuthu TR and
Iturrate I (2024), Safe contact-based robot
active search using Bayesian optimization and
control barrier functions.
Front. Robot. AI 11:1344367.
doi: 10.3389/frobt.2024.1344367

COPYRIGHT

© 2024 Vinter-Hviid, Sloth, Savarimuthu and
Iturrate. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Safe contact-based robot active search using Bayesian optimization and control barrier functions

Frederik Vinter-Hviid^{1*}, Christoffer Sloth²,
Thiusius Rajeeth Savarimuthu^{1,2} and Iñigo Iturrate²

¹ROPCA ApS, Odense, Denmark, ²SDU Robotics, Maersk McKinney Moller Institute, University of Southern Denmark, Odense, Denmark

In robotics, active exploration and learning in uncertain environments must take into account safety, as the robot may otherwise damage itself or its surroundings. This paper presents a method for safe active search using Bayesian optimization and control barrier functions. As robot paths undertaken during sampling are continuous, we consider an informative continuous expected improvement acquisition function. To safely bound the contact forces between the robot and its surroundings, we leverage exponential control barrier functions, utilizing the derivative of the force in the contact model to increase robustness to uncertainty in the contact boundary. Our approach is demonstrated on a fully autonomous robot for ultrasound scanning of rheumatoid arthritis (RA). Here, active search is a critical component of ensuring high image quality. Furthermore, bounded contact forces between the ultrasound probe and the patient ensure patient safety and better scan quality. To the best of our knowledge, our results are both the first demonstration of safe active search on a fully autonomous robot for ultrasound scanning of rheumatoid arthritis and the first experimental evaluation of bounding contact forces in the context of medical robotics using control barrier functions. The results show that when search time is limited to less than 60 s, informative continuous expected improvement leads to a 92% success, a 13% improvement compared to expected improvement. Meanwhile, exponential control barrier functions can limit the force applied by the robot to under 5 N, even in cases where the contact boundary is specified incorrectly by -1 or $+4$ mm.

KEYWORDS

control barrier function, Bayesian optimization, active search, autonomous ultrasound scanning, robot force control

1 Introduction

Active learning, along with related problems such as active object exploration and localization, find widespread applications across various robotic domains. From mobile robot navigation (Lluvia et al., 2021) to grasping (Kroemer et al., 2010) to object and scene reconstruction (Jamali et al., 2016), many robotic problems involve variations of the exploration-exploitation dilemma, where an autonomous system must simultaneously regress and optimize an unknown or uncertain function.

Several methods have been proposed to tackle the above problems, including many based on deep learning (Ren et al., 2021). For the low-data regime, which is often the case for online learning in robotics, the most common are methods based on Gaussian process regression (Jamali et al., 2016; Yi et al., 2016; Driess et al., 2019; De Farias et al., 2021) and Bayesian optimization (Nogueira et al., 2016; Goel et al., 2022; Roveda et al., 2022; Berkenkamp et al., 2023). Koopman operators have also been suggested for the case of learning the nonlinear dynamics of a physical system for control (Abraham and Murphey, 2019), but this method is only applicable to learning dynamics and not to a more general function approximation problem.

As active exploration is, by definition, an uncertain process, a key consideration is the assurance of safety despite operating in fully or partially unknown environments. Safety can be handled algorithmically in the learning process by considering, for example, the uncertainty in a Gaussian process regression (Turchetta et al., 2019). However, here we are more interested in methods that consider the safety of the controlled system, that is, the robot, as a query deemed safe algorithmically can still result in an unsafe robot action due to uncertainty in the control and environmental dynamics. Various methods have been proposed, such as safe control based on energy functions (Pandya and Liu, 2022), safe model predictive control (Koller et al., 2018), and Bayesian meta-learning through alternating sequential optimal control problems for exploration and exploitation (Lew et al., 2022). A prominent approach to safety is control barrier functions (CBFs), a framework for ensuring the safety of nonlinear control affine systems (Ames et al., 2014; Romdlony and Jayawardhana, 2014; Ames et al., 2019). The resulting equations can be solved using quadratic programming, making the implementation feasible in control loops. The methodology has been expanded to high relative degree systems, making it generally applicable (Nguyen and Sreenath, 2016). Model uncertainty has been considered in CBFs in terms of parametric uncertainty (Cohen and Belta, 2022) and non-parametric uncertainty (Castañeda et al., 2021). These works ensure that safety can be guaranteed despite having a bounded model uncertainty.

In this paper, we consider the combination of the above two problems, namely, safe active exploration, in the context of an autonomous robot system for ultrasound (US) scanning of rheumatoid arthritis (RA). US imaging is a popular method for diagnosis and monitoring of RA and various other diseases, as it allows inspection of tissue and joint structures at low cost and without the use of radiation. The use of ultrasound does, however, rely on trained professionals with limited availability, so automation of the process has high potential value.

Active exploration is a major component of US scanning for RA and is, in fact, how the procedure is performed manually by trained professionals. During the procedure, the probe is placed in contact with joints in the patient's hand, allowing the subsurface structures to be inspected on a connected monitor. The professional must ensure that contact between the probe and the patient's hand is maintained and find an optimal placement in order to capture potential disease activity. When an appropriate area has been found, the probe is kept stationary, and Doppler mode imaging is used to assess the flow of synovial fluid in the joint cavity, indicating

inflammation. Correct placement of the probe will ensure high-quality scans. Different approaches have been proposed to optimize the ultrasound image quality. A commonly used method is based on ultrasound confidence maps (Karamalis et al., 2012; Chatelain et al., 2017; Jiang et al., 2020) that can be used to detect if contact is poor and adjustments to the probe position must be made. Force-based methods have also been explored to optimize the image quality by attempting to estimate the surface normal and position of the probe accordingly (Jiang et al., 2021a). As these methods do not consider the structure of tissues in the image but only the overall quality, they are not suited for RA, where the structure is vital. Bayesian optimization based on segmented ultrasound images is proposed by Goel et al. (2022). This allows certain structures in the image to be optimized efficiently, and as the optimization is guided by a statistical model, uncertainties are handled implicitly. While the above methods consider the positioning of the probe, limiting the maximum applied force is also critical. First, the contact force must be kept low to ensure the flow of synovial fluid is not blocked by the pressure (Möller et al., 2017). Second, patients with high levels of disease activity will experience pain from excessive pressure.

To address the above use case, we propose to draw from the wider active exploration and safe control literature. To determine the probing location, we use Bayesian optimization, similar to Goel et al. (2022), but make use of a classifier trained to estimate the US image quality instead of a segmentation model. This is done because assessing the quality of an ultrasound image for RA is a complex task requiring the correct placement, appearance, and relative size of different tissue types. We extend the method by considering the ability to sample images continuously while positioning the probe on the patient, leading to a more efficient search. To establish and maintain contact with the patient's hand, we use direct force control and add an additional level of safety in the form of a high-order CBF to bound the contact forces below a safe threshold.

Our main contributions are as follows:

- We present, to the best of our knowledge, the first results on CBFs applied to safe force control in the medical domain of a real (not simulated) robotic system in a laboratory setting.
- We present a method for autonomous robot scanning of RA. While research has been performed on autonomous US scanning for various applications such as scanning of vessels (Virga et al., 2016; Jiang et al., 2021b) and breasts (Welleweerd et al., 2020), to the best of our knowledge, no previous work within autonomous scanning for RA exists.
- Additionally, a minor contribution is the extension of informative continuous expected improvement (ICEI) as an acquisition function for Bayesian optimization to the context of robot motion planning.

2 Problem formulation

We consider two sub-problems in the context of ultrasound scanning for RA:

Problem 1: (Active search). This problem is concerned with finding areas of high ultrasound image quality in the presence of uncertain measurements. We consider the search problem as finding a global maximum p_{opt} of an unknown underlying function $f(p)$ with

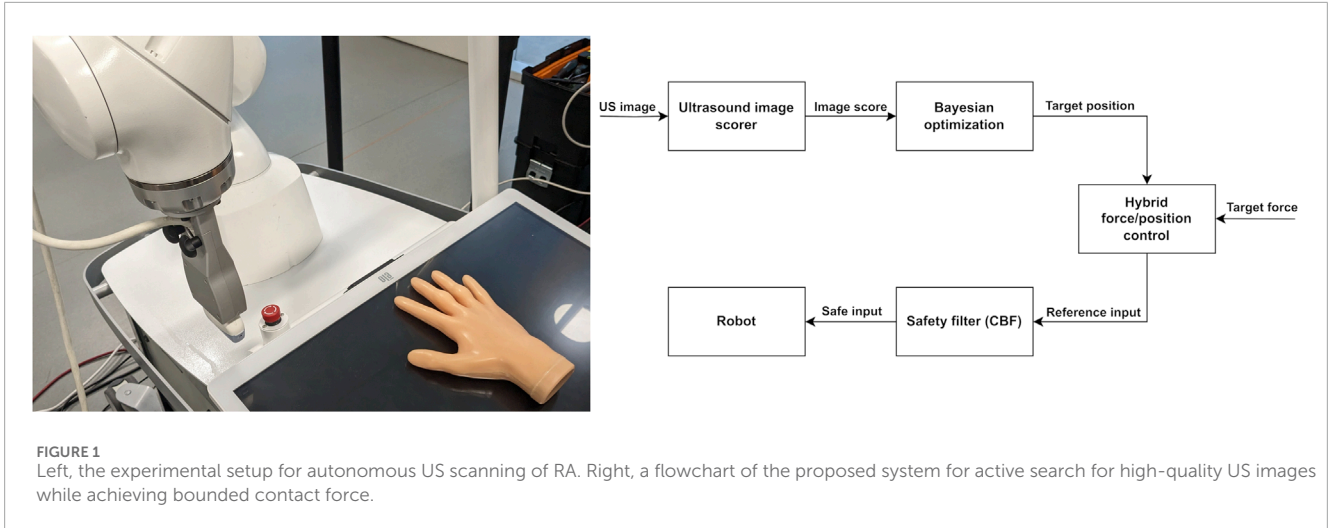


FIGURE 1 Left, the experimental setup for autonomous US scanning of RA. Right, a flowchart of the proposed system for active search for high-quality US images while achieving bounded contact force.

uncertainty ϵ where \mathbf{p} is a subset of the Cartesian space.

$$\mathbf{p}_{\text{opt}} = \arg \max_{\mathbf{p} \in \mathbb{R}^n} f(\mathbf{p}) + \epsilon. \quad (1)$$

Function estimates are achieved using a classifier trained on ultrasound images. It is assumed that a starting point is given from external sensing, such as a RGB-D camera combined with a method of estimating joint locations in a RGB image. Solutions for this exist, at least for simple, well-lit scenes (Zhang et al., 2020).

Problem 2: (Bounding the contact force). This problem consists of ensuring bounded contact force with the patient despite uncertainties in the patient’s position. A robot manipulator with an ultrasound probe mounted—as seen in Figure 1—is employed. The robot is controlled at the joint position level. We consider the problem of finding a safe control input u^* given a potentially unsafe nominal control input u_{no} that minimizes the differences between u^* and u_{no} . Safety is defined in terms of bounded contact force $f_c \leq f_{\text{max}}$, where f_c and f_{max} are the contact force and maximum allowed force, respectively.

3 Preliminaries

This section will present theoretical background on Gaussian process regression and Bayesian optimization, which are used as the basis of our solution to Problem 1, and exponential control barrier functions, which are used as part of our solution to Problem 2.

3.1 Gaussian process regression

Gaussian process regression (GPR) allows unknown functions, along with their uncertainty, to be estimated based on limited samples (Rasmussen and Williams, 2005). A Gaussian process (GP) is a distribution over functions and is defined fully by a mean function $m(\mathbf{p})$ and a covariance function $k(\mathbf{p}, \mathbf{p}')$. It is assumed that the vector input $\mathbf{p} \in \mathbb{R}^d$ is related to the scalar output y by the function $f(\cdot)$ with the following relationship: $y = f(\mathbf{p}) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. The training input is defined as

\mathbf{P} with \mathbf{y} being the associated targets. Together, this is denoted as $\mathcal{D}:\{(\mathbf{p}_1, y_1), (\mathbf{p}_2, y_2), \dots, (\mathbf{p}_n, y_n)\}$. Observations \mathbf{y} and discrete function values to be predicted f_* then follow a multivariate normal distribution:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} m(\mathbf{P}) \\ m(\mathbf{P}_*) \end{bmatrix}, \begin{bmatrix} \mathbf{K}_\epsilon & K(\mathbf{P}, \mathbf{P}_*) \\ k(\mathbf{P}_*, \mathbf{P}) & K(\mathbf{P}_*, \mathbf{P}_*) \end{bmatrix} \right), \quad (2)$$

where $\mathbf{K}_\epsilon = K(\mathbf{P}, \mathbf{P}) + \sigma_\epsilon^2 \mathbf{I}$, with σ_ϵ^2 being the inherent noise on the measurements—allowing optimization under uncertainty—and

$$K(\mathbf{P}, \mathbf{P}') = \begin{bmatrix} k(\mathbf{p}_1, \mathbf{p}'_1) & k(\mathbf{p}_1, \mathbf{p}'_2) & \dots & k(\mathbf{p}_1, \mathbf{p}'_k) \\ k(\mathbf{p}_2, \mathbf{p}'_1) & k(\mathbf{p}_2, \mathbf{p}'_2) & \dots & k(\mathbf{p}_2, \mathbf{p}'_k) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{p}_k, \mathbf{p}'_1) & k(\mathbf{p}_k, \mathbf{p}'_2) & \dots & k(\mathbf{p}_k, \mathbf{p}'_k) \end{bmatrix}. \quad (3)$$

The mean vector associated with \mathbf{f}_* is then

$$\bar{\mathbf{f}}_* = m(\mathbf{P}_*) + K(\mathbf{P}_*, \mathbf{P}) \mathbf{K}_\epsilon^{-1} (\mathbf{y} - m(\mathbf{P})) \quad (4)$$

and the covariance is given by:

$$\text{cov}(\mathbf{f}_*) = K(\mathbf{P}_*, \mathbf{P}_*) - K(\mathbf{P}_*, \mathbf{P}) \mathbf{K}_\epsilon^{-1} K(\mathbf{P}, \mathbf{P}_*). \quad (5)$$

The estimate of the standard deviation along the estimated function is then the diagonal of $\text{cov}(\mathbf{f}_*)$ and is denoted $\sigma_f(\mathbf{p})$. The mean function is often set to zero unless some underlying information about the process exists. The covariance function is generally defined such that points near each other are more correlated. A common choice is the Matern class of covariance functions. This class of functions leads to different levels of differentiability and, thus, smoothness of $f(\mathbf{p})$ based on a smoothness parameter ν . For $\nu = \frac{3}{2}$, the Matern kernel is

$$k_{\nu=\frac{3}{2}}(\mathbf{p}, \mathbf{p}') = (1 + \sqrt{3}d(\mathbf{p}, \mathbf{p}')) \exp(-\sqrt{3}d(\mathbf{p}, \mathbf{p}')), \quad (6)$$

where $d(\mathbf{p}, \mathbf{p}')$ is the Euclidean distance between \mathbf{p} and \mathbf{p}' scaled by a diagonal matrix Θ :

$$d = (\mathbf{p} - \mathbf{p}') \Theta^{-1} (\mathbf{p} - \mathbf{p}'), \quad (7)$$

$$\Theta = \text{diag}(l_1, l_2, \dots, l_d) = \text{diag}(\theta), \quad (8)$$

with $\theta = [l_1 \ l_2 \ \dots \ l_d]^T$ as the length scales of the function. The length scales encode how rapidly $f(\cdot)$ changes along each dimension. These are considered free parameters that must be determined from the training data. This is done by maximizing the marginal log-likelihood of the training examples given the parameters.

3.2 Bayesian optimization

Bayesian optimization provides a framework for maximizing functions that are costly to evaluate, where uncertainty is present, and where following the gradient is infeasible (Shahriari et al., 2016). The method consists of three elements: a statistical surrogate model, an acquisition function, and a stopping criterion. A common choice for the surrogate model is Gaussian processes, as defined above. The acquisition function $\alpha(\cdot)$ determines where to sample next in the search space based on the surrogate model by balancing exploration and exploitation:

$$\mathbf{p}_{n+1} = \arg \max_{\mathbf{p} \in \mathbb{R}^d} \alpha(\mathbf{p}; \mathcal{D}). \quad (9)$$

A common choice for the acquisition function is expected improvement (EI) (Mockus et al., 1978). EI maximizes the expected gain over the previous best sampled score, f^* .

$$EI(\mathbf{p}) := \mathbb{E}[\max((f(\mathbf{p}) - f^*), 0)], \quad (10)$$

EI can be calculated analytically as

$$EI(\mathbf{p}) = \begin{cases} (\bar{f}_*(\mathbf{p}) - f^*) \Phi\left(\frac{\bar{f}_*(\mathbf{p}) - f^*}{\sigma_f(\mathbf{p})}\right) + \sigma_f(\mathbf{p}) \phi\left(\frac{\bar{f}_*(\mathbf{p}) - f^*}{\sigma_f(\mathbf{p})}\right) & \text{if } \sigma_f(\mathbf{p}) > 0 \\ 0 & \text{if } \sigma_f(\mathbf{p}) = 0 \end{cases}, \quad (11)$$

where $\Phi(\cdot)$, $\phi(\cdot)$ are the cumulative distribution function (CDF) and the probability density function (PDF), respectively.

3.3 Control barrier functions

Control barrier functions provide a framework for ensuring system safety (Ames et al., 2014). Throughout this section, a nonlinear control affine system of the following form is considered:

$$\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}, \quad (12)$$

where $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are local Lipschitz functions, with $\mathbf{x} \in X \subset \mathbb{R}^n$ and $\mathbf{u} \in U \subset \mathbb{R}^m$ being the admissible states and control inputs, respectively. The goal of the CBF is to keep the system in a safe region of the state space. This safe region \mathcal{C} is defined based on a continuously differentiable function $h: X \subset \mathbb{R}^n \rightarrow \mathbb{R}$ as

$$\begin{aligned} \mathcal{C} &= \{\mathbf{x} \in X \mid h(\mathbf{x}) \geq 0\}, \\ \partial\mathcal{C} &= \{\mathbf{x} \in X \mid h(\mathbf{x}) = 0\}, \\ \text{Int}(\mathcal{C}) &= \{\mathbf{x} \in X \mid h(\mathbf{x}) > 0\}. \end{aligned} \quad (13)$$

The system is thus safe if a control law input \mathbf{u} leads to the scalar function $h(\mathbf{x}) \geq 0 \ \forall t \wedge \forall \mathbf{x}_0 \in \mathcal{C}$. Equivalently, this can be described as

the set \mathcal{C} being rendered forward invariant by the control law. When the relative degree r of the system with respect to h is larger than one, higher-order CBFs can be used to render a set \mathcal{C} safe. The relative degree of a continuously differentiable function $h(\mathbf{x})$ with respect to System (12) is the number of times it can be differentiated along (12) before the control input \mathbf{u} explicitly shows. Exponential control barrier functions (ECBFs) are a class of higher-order CBFs (Nguyen and Sreenath, 2016). This formulation is based on calculating higher-order time derivatives of $h(\mathbf{x})$. The r^{th} time derivative is

$$h^{(r)}(\mathbf{x}, \mathbf{u}) = L_f^r h(\mathbf{x}) + L_g L_f^{r-1} h(\mathbf{x}) \mathbf{u}, \quad (14)$$

where $L_f^r h(\mathbf{x})$ is the r^{th} Lie derivative of h along f . ECBFs are defined based on a series of integrators that relate $h^{(r)}(\mathbf{x}, \mathbf{u})$ to $h(\mathbf{x})$:

$$\begin{aligned} \dot{\eta}(\mathbf{x}) &= \mathbf{F}\eta(\mathbf{x}) + \mathbf{G}\mu, \\ h(\mathbf{x}) &= \mathbf{C}\eta(\mathbf{x}). \end{aligned} \quad (15)$$

With

$$\begin{aligned} \eta(\mathbf{x}) &= \begin{bmatrix} h(\mathbf{x}) \\ L_f h(\mathbf{x}) \\ L_f^2 h(\mathbf{x}) \\ \vdots \\ L_f^{r-1} h(\mathbf{x}) \end{bmatrix} = \begin{bmatrix} h(\mathbf{x}) \\ \dot{h}(\mathbf{x}) \\ \ddot{h}(\mathbf{x}) \\ \vdots \\ h^{(r-1)}(\mathbf{x}) \end{bmatrix}, \\ \mu &= L_f^r h(\mathbf{x}) + L_g L_f^{r-1} h(\mathbf{x}) \mathbf{u} = h^{(r)}(\mathbf{x}, \mathbf{u}) \end{aligned} \quad (16)$$

and $\mathbf{F}: r \times r$, $\mathbf{G}: r \times 1$, $\mathbf{C}: 1 \times r$

$$\mathbf{F} = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}, \quad \mathbf{C} = [1 \ 0 \ \dots \ 0.] \quad (17)$$

Choosing the state feedback $\mu = -\mathbf{K}\eta(\mathbf{x})$ leads to $h(\mathbf{x})$ being an explicit function of time t and the initial state \mathbf{x}_0 :

$$h(\mathbf{x}) = \mathbf{C}e^{(\mathbf{F}-\mathbf{GK})t}\eta(\mathbf{x}_0), \quad (18)$$

where $h(\mathbf{x})$ is then said to be an ECBF—rendering \mathcal{C} forward invariant—if there exists a row vector \mathbf{K}_y such that

$$\sup_{\mathbf{u} \in U} [L_f^r h(\mathbf{x}) + L_g L_f^{r-1} h(\mathbf{x}) \mathbf{u}] \geq -\mathbf{K}_y \eta(\mathbf{x}), \ \forall \mathbf{x} \in \text{Int}(\mathcal{C}) \quad (19)$$

results in

$$h(\mathbf{x}(t)) \geq \mathbf{C}e^{-\mathbf{GK}_y t} \eta(\mathbf{x}_0) \geq 0, \quad (20)$$

where $h(\mathbf{x}_0) \geq 0$. A \mathbf{K}_y can be found that satisfies this using the pole placement method, with all poles being real and negative. Additionally, the eigenvalues λ_i of system $\mathbf{F} - \mathbf{GK}_y$ must adhere to a condition based on the initial conditions \mathbf{x}_0 :

$$\lambda_i(\mathbf{F} - \mathbf{GK}_y) \geq -\frac{\dot{v}_{i-1}(\mathbf{x}_0)}{v_{i-1}(\mathbf{x}_0)}, \quad (21)$$

where

$$\begin{aligned} v_0(\mathbf{x}) &= h(\mathbf{x}), \\ v_1(\mathbf{x}) &= \dot{v}_0(\mathbf{x}) + s_1 v_0(\mathbf{x}), \\ &\vdots \\ v_r(\mathbf{x}) &= \dot{v}_{r-1}(\mathbf{x}) + s_r v_{r-1}(\mathbf{x}), \end{aligned} \quad (22)$$

and s_0, \dots, s_r are defined as the roots of the characteristic polynomial of $\mathbf{F} - \mathbf{G}\mathbf{K}_y$.

4 Approach

This section presents our proposed method for safe active search under uncertainty applied to ultrasound image acquisition.

4.1 Ultrasound image quality optimization

We will now present our proposed solution to Problem 1. In regular Bayesian optimization, sampling from anywhere in the search space is assumed to be associated with the same cost. As the robot must move to acquire images, this assumption is not ideal for robotic ultrasound scanning, as switching between sampling points far apart from each other will incur a large time cost while the robot moves between them. To account for this, we consider a non-standard acquisition function that incorporates the ability to sample ultrasound images throughout the path traveled by the robot. Similar methods have been applied to robotic environmental monitoring using unmanned aerial vehicles (Marchant and Ramos, 2014). We parameterize the path as straight lines to ensure the optimization time of the acquisition function is kept low. The path \mathcal{P} is defined as

$$\mathcal{P}(u, \mathbf{p}_{n+1} | \mathbf{p}_n) = (1 - u)\mathbf{p}_{n+1} + u\mathbf{p}_n, \quad (23)$$

with $u \in [0, 1]$. The modified acquisition function is denoted $\alpha_{IC}(\cdot)$ and is said to be informative continuous as information throughout the path is considered. It is defined by integrating the base acquisition function over the path and scaling with the inverse of the path length:

$$\alpha_{IC}(\mathcal{P}(u, \mathbf{p}_{n+1} | \mathbf{p}_n), \alpha) = \frac{1}{\|\mathbf{p}_{n+1} - \mathbf{p}_n\|_2} \int_0^1 \alpha(\mathcal{P}(u, \mathbf{p}_{n+1} | \mathbf{p}_n)) du. \quad (24)$$

The scale is included to account for the time spent moving between \mathbf{p}_n and \mathbf{p}_{n+1} . Assuming that the robot moves at a constant velocity, maximizing the modified acquisition function leads to a path that maximizes the integral of the base acquisition function per unit of time spent. The formulation leads to the entire path being considered—and thus less time spent in areas with low uncertainty and/or low expected scores. We employ the EI acquisition function as the base acquisition function and denote the informative continuous version as ICEI. A score y is associated with each position in the search space \mathbf{p} using a classifier trained on labeled ultrasound images. The range of the output is $[0; 1]$, with 0 indicating poor quality and 1 good. Defining an appropriate stopping criterion is important to ensure a good area has been found without spending excessive time scanning a patient. A statistical stopping criterion is proposed based on the maximum lower bound over the samples, \mathbf{y} . The lower bound for a sample is

$$y_{1-\delta} = y - Z_\delta \sigma_y, \quad (25)$$

where Z is the standard normal distribution, Z_δ is the critical value for a confidence level of $1 - \delta$, and σ_y is the prediction of standard

deviation for the corresponding sample y based on the surrogate model. As the range of possible scores is known beforehand, defining an appropriate threshold is possible.

4.2 Safe direct force control

Our proposed solution to Problem 2 combines a hybrid force/position controller—also referred to as the nominal controller—with a CBF defined such that bounded contact force is ensured. The hybrid force/position controller allows direct control over the force to ensure consistent ultrasound image quality while making positioning possible during contact. The robot considered in this work is controlled at the joint position level. The CBF is defined in Cartesian space and at the acceleration level to allow for constraints on the force. End-effector accelerations are related to joint accelerations using the following kinematic relationship:

$$\ddot{\mathbf{x}} = \mathbf{J}^\dagger(\mathbf{q})(\ddot{\mathbf{x}} - \dot{\mathbf{J}}(\mathbf{q})\dot{\mathbf{q}}), \quad (26)$$

where $\mathbf{J}(\mathbf{q})$ is the manipulator Jacobian, † denotes the Moore–Penrose pseudo inverse, \mathbf{q} is the joint configuration, and $\ddot{\mathbf{x}}$ is the Cartesian space acceleration. The resulting joint accelerations are then double integrated for the robot’s position controller. Figure 2 illustrates the proposed control architecture.

4.2.1 Hybrid force/position controller

The error between the desired and measured end-effector wrench is denoted $\Delta\mathbf{h}$. Similarly, the positional error used is denoted $\Delta\mathbf{x}$; note that the orientational error is obtained using quaternions. The Cartesian space is divided such that force control is applied in some subset of the space, and position control is applied in the remainder of the space (Siciliano et al., 2008). New variables are defined to achieve this:

$$\Delta\mathbf{h}_{m \times 1}^* = \mathbf{S}_f^\dagger \frac{\Delta\mathbf{h}}{m \times 6}, \quad (27)$$

$$\Delta\mathbf{x}_{n \times 1}^* = \mathbf{S}_p^\dagger \frac{\Delta\mathbf{x}}{n \times 6}, \quad (28)$$

where \mathbf{S}_p^\dagger and \mathbf{S}_f^\dagger select the position-controlled and force-controlled subsets, respectively. Force control is applied along the length of the probe, defined as the z -axis, while positional control is used for the remaining degrees of freedom (DOFs). The force controller acting on a subset of the Cartesian space is then

$$\bar{\mathbf{a}}_f = \mathbf{K}_p^f \mathbf{S}_f^\dagger \Delta\mathbf{h} - \mathbf{K}_d^f \mathbf{S}_f^\dagger \dot{\mathbf{x}}_e, \quad (29)$$

where $\dot{\mathbf{x}}_e$ is the end-effector velocity, $\bar{\mathbf{a}}_f$ is the resulting acceleration reference, and $\mathbf{K}_p^f, \mathbf{K}_d^f$ are diagonal $m \times m$ positive definite gain matrices. Note that $\dot{\mathbf{x}}_e$ is used over the more natural choice of using the end-effector wrench derivative $\dot{\mathbf{h}}_e$. This is a method of avoiding the problem of taking the time derivative of the often highly noisy force measurements (Siciliano et al., 2008). The position controller is defined similarly as

$$\bar{\mathbf{a}}_p = \mathbf{K}_p^p \mathbf{S}_p^\dagger \Delta\mathbf{x} - \mathbf{K}_d^p \mathbf{S}_p^\dagger \dot{\mathbf{x}}_e, \quad (30)$$

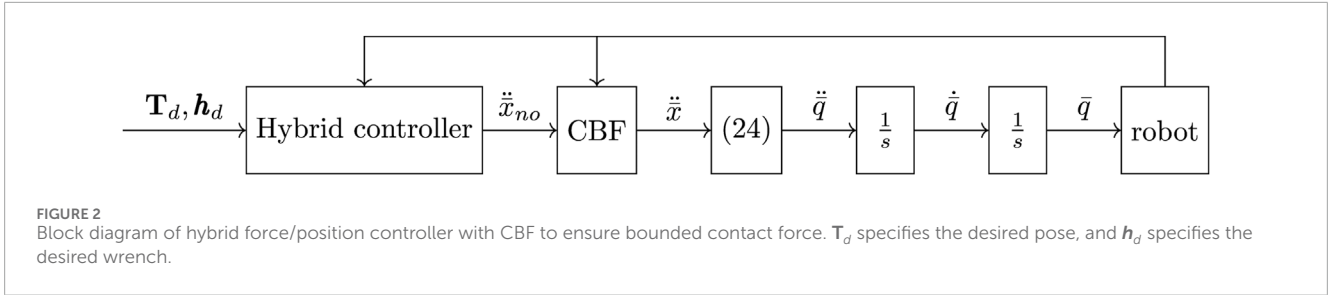


FIGURE 2 Block diagram of hybrid force/position controller with CBF to ensure bounded contact force. T_d specifies the desired pose, and h_d specifies the desired wrench.

with \bar{a}_p as the acceleration reference, and K_p^p, K_d^p are $n \times n$ positive definite gain matrices. The full acceleration of the nominal controller is then

$$\ddot{x}_{no} = S_p \bar{a}_f + S_p \bar{a}_p. \tag{31}$$

4.2.2 CBF for bounded contact force

The acceleration reference from the hybrid force/position control is modified by a control barrier function to ensure bounded contact force. This requires the definition of a control affine model relating commanded end-effector accelerations to contact force. This is defined as being one-dimensional, only considering the force from movement along the z -axis, thus neglecting friction in the axes of motion. This is a valid assumption, as the US gel placed on the patient’s hand while scanning will render friction between the probe and hand negligible. The contact force is modeled using the Kelvin–Voigt model as

$$f_c = k(z - z_r) + b\dot{z}, \tag{32}$$

where f_c is the contact force, k and b are positive scalars, and z_r is the position of the contact surface.

A second-order system is used to model the inner position control loop and account for the actuator dynamics. The relationship between position reference \bar{z} and actual position z is modeled by a second-order system:

$$\ddot{z} = \omega_n^2(\bar{z} - z) + 2\zeta\omega_n\dot{z}, \tag{33}$$

where ζ is the damping factor and ω_n is the undamped natural frequency. The control input to the CBF from the nominal force controller is defined at the acceleration level $u = \ddot{\bar{z}}$. To relate this to \bar{z} , it is double integrated:

$$\bar{Z}(s) = \frac{1}{s^2} \ddot{Z}(s). \tag{34}$$

The full control affine model is then

$$\mathbf{x} = \begin{bmatrix} f_c \\ z \\ \dot{z} \\ \bar{z} \\ \dot{\bar{z}} \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}, \quad u = \ddot{\bar{z}} \tag{35}$$

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \dot{x}_4 \\ \dot{x}_5 \end{bmatrix} = \begin{bmatrix} kx_3 + b(\omega_n^2(x_4 - x_2) + 2\zeta\omega_n x_3) \\ x_3 \\ \omega_n^2(x_4 - x_2) + 2\zeta\omega_n x_3 \\ x_5 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} u.$$

Note that this utilizes the derivative of the force, such that the method does not directly depend on z_r . However, it will lead to an overly conservative barrier function outside of contact, as the non-contact situation is not modeled. The barrier function is defined as

$$h(\mathbf{x}) = f_{max} - f_c, \tag{36}$$

where f_{max} is the maximum allowable force. This leads to a relative degree of three for the system, and an ECBF is therefore employed. The ECBF is defined as a constraint to a quadratic program (QP). The objective function is defined as the squared error between the nominal control input u_{no} and the resulting input u

$$|u - u_{no}|^2 = u^2 - 2u_{no}u + u_{no}^2. \tag{37}$$

This ensures the nominal control is followed whenever possible, and the ECBF only limits the input whenever the system is near the boundary of the safe set. The resulting QP is then

$$u^*(\mathbf{x}, u_{no}) = \underset{u \in \mathbb{R}}{\operatorname{argmin}} \quad u^2 - 2u_{no}u \tag{38}$$

s.t. $L_f^r h(\mathbf{x}) + L_g L_f^{r-1} h(\mathbf{x}) u + K_y \eta(\mathbf{x}) \geq 0,$

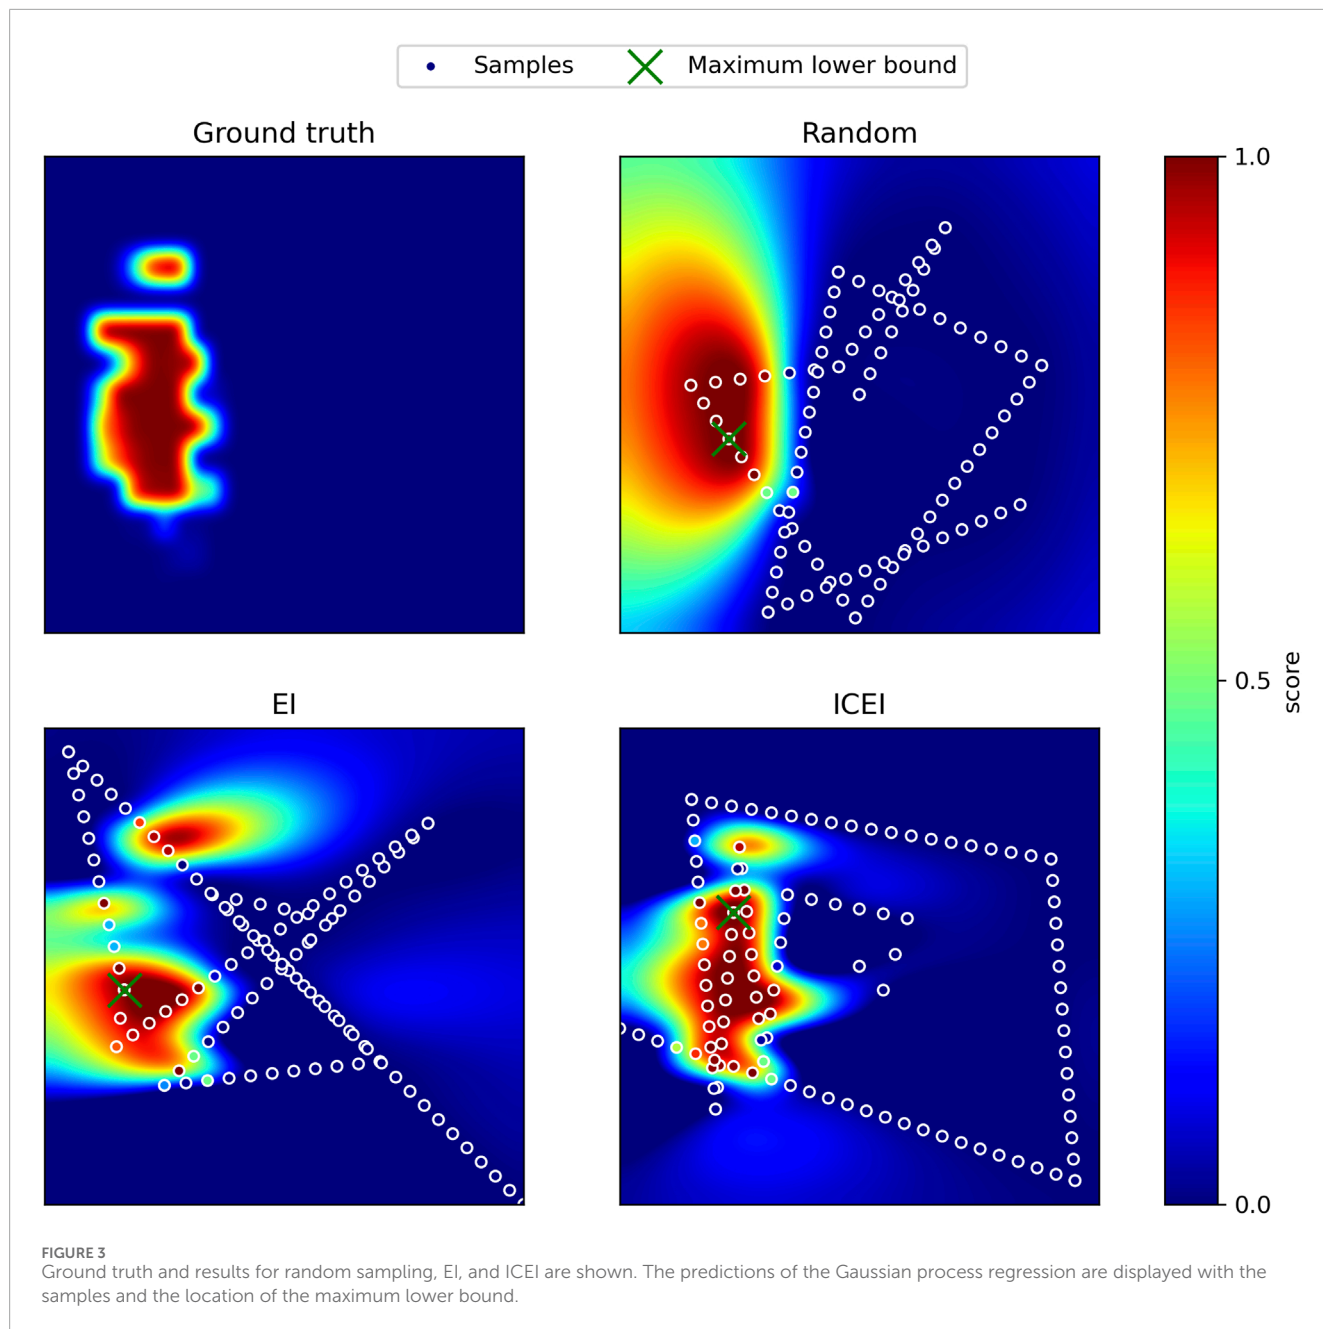
where u^* is the modified control input and u_{no} is the nominal control input from the hybrid force/position controller.

5 Results

In this section, results for the image quality optimization method and safe force controller are presented. The standard EI acquisition function is compared with the proposed ICEI variant. The addition of a CBF for keeping bounded contact force is evaluated under uncertainty in the estimate of surface location z_r .

5.1 Comparing image quality optimization methods

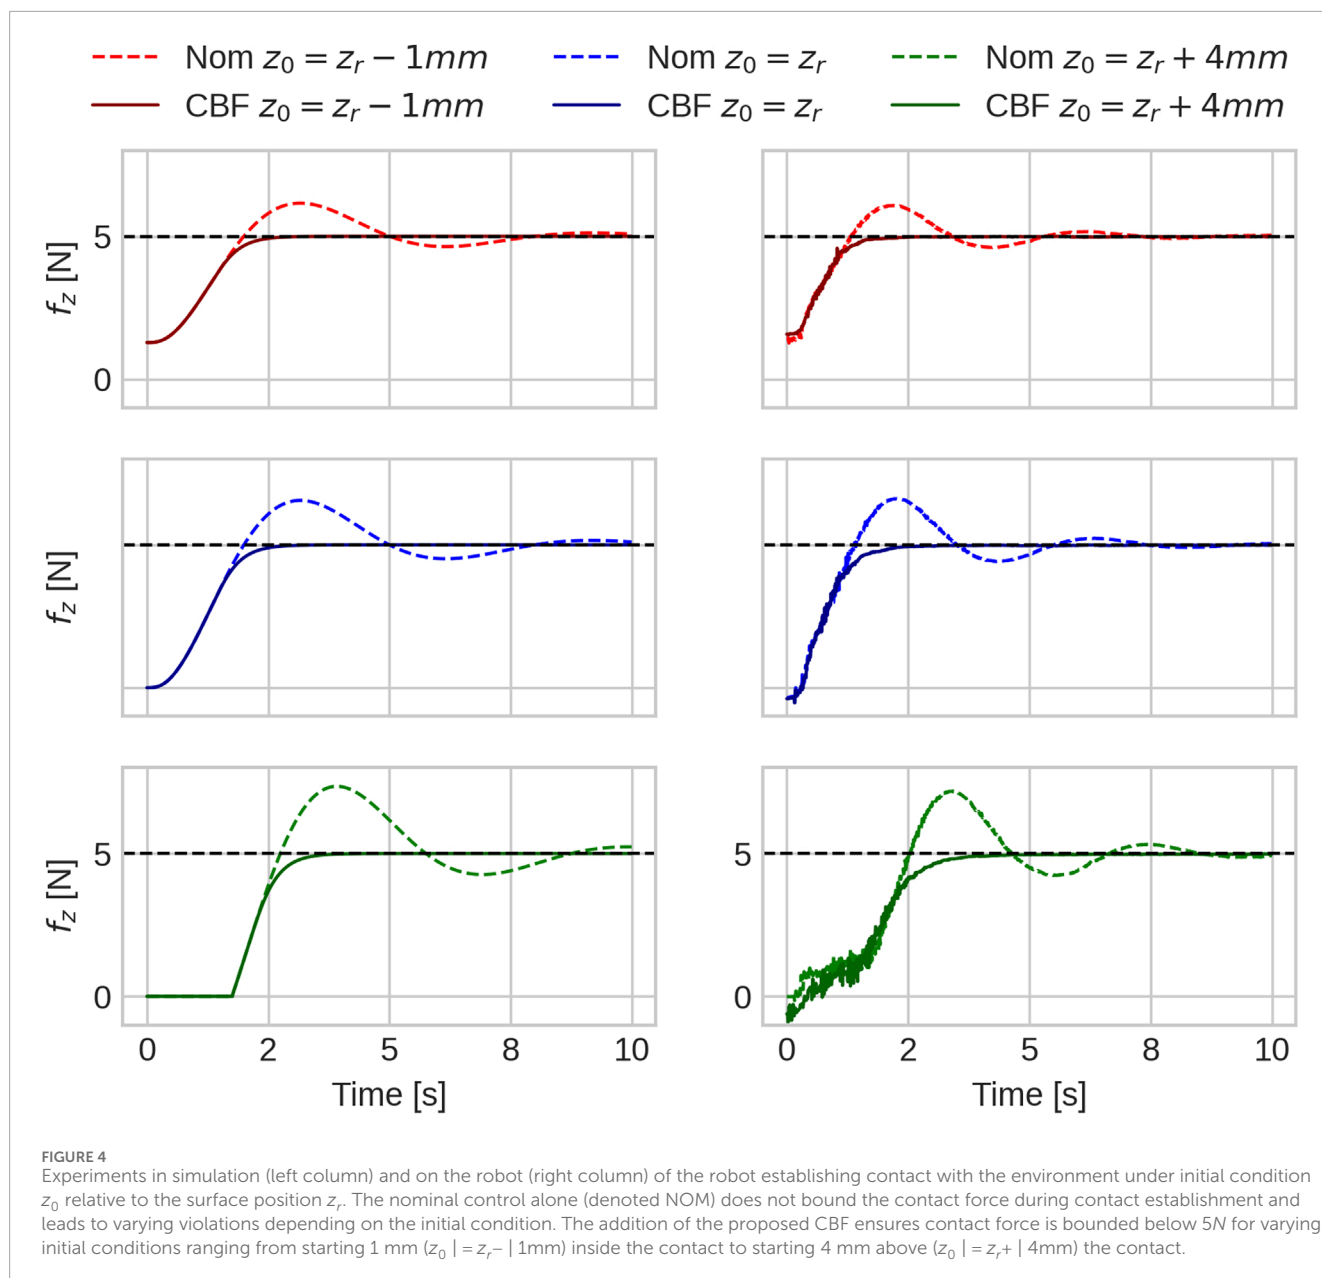
The ultrasound image quality optimization approach assumes a starting position is given, and predetermined bounds of the search space are defined in relation to this. The search was performed over two DOFs such that the input to the Bayesian optimizer was $\mathbf{p} = [x, R_x]^T$, where R_x is rotation about the x -axis. The inherent noise of the measurements was set to $\sigma_\epsilon = 0.1$. Bayesian optimization was implemented using BoTorch (Balandat et al., 2020). The informative continuous extension was implemented in PyTorch, estimating the integral using a discrete sum. Our proposed method for optimizing



ultrasound image quality has been evaluated in a simulated setting. The simulation is based on data gathered from scanning 32 joints across four test subjects. The search space was sampled in a grid using nearest-neighbor interpolation to construct an approximated ground truth. The stopping criterion was set to a maximum lower bound of 0.9 with a confidence level of 95%. Three methods were tested: random sampling, EI, and ICEI. The search space was normalized such that the range was equal for the two DOFs. A distance budget of 10 times the length of the search space was given. In order to simulate the ability to sample continuously, a sample was taken for every 4% of the search space traveled. Based on this, random sampling, EI, and ICEI resulted in 60.7%, 67.8%, and 75.0% success rates, respectively, with mean distance traveled being 8.65, 8.13 and 7.64 times the length of the normalized search

space respectively. Figure 3 shows the search patterns resulting from the different methods. Random sampling is, as expected, inefficient. The tendency of ICEI to cause less overlap of paths compared to EI is clear.

EI and ICEI were also compared in a live test. In these tests the search was extended such that $p = [x, R_x, R_y, R_z]^T$ with a range of $[-7.5mm, 7.5mm]$ for x and $[-4^\circ, 4^\circ]$ for the rotational DOFs. The test compared the results of searching using EI and ICEI. For each test, the methods were started at the same position. The methods were run until the stopping criterion was reached. When this occurred, the probe was positioned at the maximum lower bound, and a measurement was taken to ensure the success criterion actually had been reached. This cannot be guaranteed due to uncertainties in positioning. If not, the search was continued. If a



solution had not been found within 60 s, it was considered a failed scan. Fourteen joints were scanned. EI lead to a success rate of 79% and ICEI to a rate of 92%. The mean recorded time for optimizing the acquisition function during this test was 0.10 ± 0.06 s for EI and 0.14 ± 0.10 s for ICEI. The training time for the Gaussian process using maximum log-likelihood was 0.16 ± 0.08 s. Even with the ICEI acquisition method, querying a new sample point only took around 0.3 s

5.2 Evaluation of safe force controller

A KUKA AG IIWA 7 R800 was used to evaluate the safe controller with the fast research interface (FRI) for communication with the robot through the IIWA ROS package

(Chatzilygeroudis et al., 2019). The communication rate with the inner joint position controller was set to 200 Hz. The *KYOTO KAGAKU Rheumatism Hand Phantom* was used with parameters estimated to $\hat{k} = 1285$ and $\hat{b} = 15$. For the inner joint position controller, the following parameters were found: $\hat{\omega}_n = 31.25$ and $\hat{\zeta} = 1$; it is not possible to change these parameters, as it is an inbuilt controller. The parameters of the joint controller were estimated by minimizing the sum of squared errors between the model and measurements taken during an initial experiment where the robot was moved into contact with the hand phantom.

The parameter K_y used in the safety filter must be defined such that the poles of System (20) are negative and real and adhere to the constraints on initial conditions (21). Pole placement within these constraints affects how conservative the barrier function is and

will always ensure safety. The following poles \mathbf{p} were utilized in the experiments: $\mathbf{p} = [-3 \quad -6 \quad -9]$. A maximum force of $f_{\max} = 5N$ was specified.

The nominal force control parameters were set to $\mathbf{K}_p^f = 0.001$ and $\mathbf{K}_d^f = 1$. The controller was implemented in Python using CVXOPT (Andersen et al., 2020) to solve the ECBF QP problem. The end-effector wrench was estimated from the robot's joint torque sensors using the following kinematic relationship:

$$\mathbf{h}_e = (\mathbf{J}(\mathbf{q})^T)^\dagger \boldsymbol{\tau}_{\text{external}}, \quad (39)$$

where $\boldsymbol{\tau}_{\text{external}}$ are the measured external joint torques available through the FRI. Figure 4 illustrates the invariance of the proposed method to the contact surface point in the simulation and when running on the robot and shows that the ECBF is able to bound the contact forces below the specified threshold.

6 Discussion

Our methodology is applicable to other robotic applications involving safety-constrained active exploration where contact between the robot and an unknown environment is concerned, such as tactile exploration and manipulation. Our proposed formulation for search using Bayesian optimization with ICEI can help model the kind of sampling scheme applicable to robot tactile exploration, that is, where the full travel path of the robot end-effector must be considered. Our approach to the safe bounding of the contact forces using ECBFs is applicable in the case of robots in contact with the environment where force feedback is available. By directly measuring the contact force and considering its first derivative in the contact model, we increase the system's robustness to uncertainty in the contact boundary, that is, uncertainty in the switching of the environment dynamics.

However, our method has a number of limitations in real-world applications. Although our CBF formulation is robust to uncertainty in the contact boundary, it is not robust to a time-varying boundary. In an ultrasound scanning application, this might occur if the patient, for example, lifted their hand during the scan. In this case, knowledge of the derivative of the contact location would be needed to guarantee safety. Similarly, the Bayesian optimization approach assumes a static function, that is, that the patient's hand does not move. Nevertheless, small motions can be accounted for by increasing the uncertainty term σ_c .

7 Conclusion

Methods for safe active exploration of an unknown environment have been presented in this article. Active search under uncertainty is performed using Bayesian optimization, and safety is achieved through the use of control barrier functions. The approach has been applied to autonomous ultrasound scanning of rheumatoid arthritis, which involves finding areas of high ultrasound image quality by positioning the ultrasound probe on the patient. The method has been demonstrated in both a simulated and real setting. An extension to the standard expected improvement, referred

to as informative continuous expected improvement, has been proposed. The method has been shown to increase the success rate from 79% to 92%, given a limited time budget of 60 s. Safety in the context of autonomous ultrasound scanning is defined in terms of keeping contact force bounded. We have shown both in simulation and on a robot that, by using a higher-order control barrier function, we are able to bound contact forces during contact establishment despite incorrectly specifying the contact boundary by -1 or $+4$ mm.

Data availability statement

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

Author contributions

FV-H: writing—original draft and writing—review and editing. II: writing—original draft and writing—review and editing. CS: writing—review and editing and writing—original draft. TS: writing—review and editing and writing—original draft.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. The work undertaken in this paper was supported by the Odense Robotics Control and Learning of Contact Transitions (CoLeCT) project, funded by the Danish Ministry of Higher Education and Science, and by the SDU I4.0Lab at the University of Southern Denmark.

Acknowledgments

The authors would also like to thank ROPCA ApS for allowing them to use their robot platform, ARTHUR, for data acquisition.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors, and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Abraham, I., and Murphey, T. D. (2019). Active learning of dynamics for data-driven control using koopman operators. *IEEE Trans. Robotics* 35, 1071–1083. doi:10.1109/tro.2019.2923880
- Ames, A. D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., and Tabuada, P. (2019). “Control barrier functions: theory and applications,” in 2019 18th European Control Conference (ECC), Naples, Italy, 25–28 June 2019 (IEEE), 3420–3431. doi:10.23919/ECC.2019.8796030
- Ames, A. D., Grizzle, J. W., and Tabuada, P. (2014). “Control barrier function based quadratic programs with application to adaptive cruise control,” in 53rd IEEE Conference on Decision and Control, Los Angeles, CA, USA, 15–17 December 2014 (IEEE), 6271–6278.
- Andersen, M., Dahl, J., and Vandenbergh, L. (2020). *CVXOPT: convex optimization*.
- Balandat, M., Karrer, B., Jiang, D. R., Daulton, S., Letham, B., Wilson, A. G., et al. (2020). BoTorch: a framework for efficient monte-carlo bayesian optimization. *Adv. Neural Inf. Process. Syst.* 33.
- Berkenkamp, F., Krause, A., and Schoellig, A. P. (2023). Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Mach. Learn.* 112, 3713–3747. doi:10.1007/s10994-021-06019-1
- Castañeda, F., Choi, J. J., Zhang, B., Tomlin, C. J., and Sreenath, K. (2021). “Pointwise feasibility of Gaussian process-based safety-critical control under model uncertainty,” in 2021 60th IEEE Conference on Decision and Control (CDC) (IEEE), 6762–6769.
- Chatelain, P., Krupa, A., and Navab, N. (2017). Confidence-Driven control of an ultrasound probe. *IEEE Trans. Robotics* 33, 1410–1424. doi:10.1109/TRO.2017.2723618
- Chatzilygeroudis, K., Mayr, M., Fichera, B., and Billard, A. (2019). Iiwa.os: a ROS stack for KUKA’s IIWA robots using the fast research interface
- Cohen, M. H., and Belta, C. (2022). “High order robust adaptive control barrier functions and exponentially stabilizing adaptive control lyapunov functions,” in 2022 American Control Conference (ACC), Atlanta, GA, USA, 08–10 June 2022 (IEEE), 2233–2238. doi:10.23919/ACC53348.2022.9867633
- De Farias, C., Marturi, N., Stolkin, R., and Bekiroglu, Y. (2021). Simultaneous tactile exploration and grasp refinement for unknown objects. *IEEE Robotics Automation Lett.* 6, 3349–3356. doi:10.1109/lra.2021.3063074
- Driess, D., Hennes, D., and Toussaint, M. (2019). “Active multi-contact continuous tactile exploration with Gaussian process differential entropy,” in 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019 (IEEE), 7844–7850.
- Goel, R., Abhimanyu, F., Patel, K., Galeotti, J., and Choset, H. (2022). “Autonomous ultrasound scanning using bayesian optimization and hybrid force control,” in 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23–27 May 2022 (IEEE), 8396–8402. doi:10.1109/ICRA46639.2022.9812410
- Jamali, N., Ciliberto, C., Rosasco, L., and Natale, L. (2016). “Active perception: building objects’ models using tactile exploration,” in 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids), Cancun, Mexico, 15–17 November 2016 (IEEE), 179–185.
- Jiang, Z., Grimm, M., Zhou, M., Esteban, J., Simson, W., Zahnd, G., et al. (2020). Automatic normal positioning of robotic ultrasound probe based only on confidence map optimization and force measurement. *IEEE Robotics Automation Lett.* 5, 1342–1349. doi:10.1109/LRA.2020.2967682
- Jiang, Z., Grimm, M., Zhou, M., Hu, Y., Esteban, J., and Navab, N. (2021a). Automatic force-based probe positioning for precise robotic ultrasound acquisition. *IEEE Trans. Industrial Electron.* 68, 11200–11211. doi:10.1109/TIE.2020.3036215
- Jiang, Z., Li, Z., Grimm, M., Zhou, M., Esposito, M., Wein, W., et al. (2021b). Autonomous robotic screening of tubular structures based only on real-time ultrasound imaging feedback. arXiv:2011.00099.
- Karamalis, A., Wein, W., Klein, T., and Navab, N. (2012). Ultrasound confidence maps using random walks. *Med. Image Anal.* 16, 1101–1112. doi:10.1016/j.media.2012.07.005
- Koller, T., Berkenkamp, F., Turchetta, M., and Krause, A. (2018). “Learning-based model predictive control for safe exploration,” in 2018 IEEE conference on decision and control (CDC), Miami, FL, USA, 17–19 December 2018 (IEEE), 6059–6066.
- Kroemer, O. B., Detry, R., Piater, J., and Peters, J. (2010). Combining active learning and reactive control for robot grasping. *Robotics Aut. Syst.* 58, 1105–1116. doi:10.1016/j.robot.2010.06.001
- Lew, T., Sharma, A., Harrison, J., Byland, A., and Pavone, M. (2022). Safe active dynamics learning and control: a sequential exploration–exploitation framework. *IEEE Trans. Robotics* 38, 2888–2907. doi:10.1109/tro.2022.3154715
- Lluvia, I., Lazkano, E., and Ansuategi, A. (2021). Active mapping and robot exploration: a survey. *Sensors* 21, 2445. doi:10.3390/s21072445
- Marchant, R., and Ramos, F. (2014). “Bayesian Optimisation for informative continuous path planning,” in 2014 IEEE International Conference on Robotics and Automation (ICRA), Hong Kong, China, 31 May 2014 - 07 June 2014 (IEEE), 6136–6143. doi:10.1109/ICRA.2014.6907763
- Mockus, J., Tiesis, V., and Zilinskas, A. (1978). The application of Bayesian methods for seeking the extremum. *Towards Glob. Optim.* 2, 2.
- Möller, I., Janta, I., Backhaus, M., Ohrndorf, S., Bong, D. A., Martinoli, C., et al. (2017). The 2017 EULAR standardised procedures for ultrasound imaging in rheumatology. *Ann. Rheumatic Dis.* 76, 1974–1979. doi:10.1136/annrheumdis-2017-211585
- Nguyen, Q., and Sreenath, K. (2016). “Exponential Control Barrier Functions for enforcing high relative-degree safety-critical constraints,” in 2016 American Control Conference (ACC), Boston, MA, USA, 06–08 July 2016 (IEEE), 322–328. doi:10.1109/ACC.2016.7524935
- Nogueira, J., Martinez-Cantin, R., Bernardino, A., and Jamone, L. (2016). “Unscented bayesian optimization for safe robot grasping,” in 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Daejeon, Korea, 09–14 October 2016 (IEEE), 1967–1972.
- Pandya, R., and Liu, C. (2022). “Safe and efficient exploration of human models during human-robot interaction,” in 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022 (IEEE), 6708–6715.
- Rasmussen, C. E., and Williams, C. K. I. (2005). *Gaussian processes for machine learning*. United States: The MIT Press. doi:10.7551/mitpress/3206.001.0001
- Ren, P., Xiao, Y., Chang, X., Huang, P.-Y., Li, Z., Gupta, B. B., et al. (2021). A survey of deep active learning. *ACM Comput. Surv. (CSUR)* 54, 1–40. doi:10.1145/3472291
- Romdlony, M. Z., and Jayawardhana, B. (2014). “Uniting control lyapunov and control barrier functions,” in 53rd IEEE Conference on Decision and Control, Los Angeles, CA, USA, 15–17 December 2014 (IEEE), 2293–2298. doi:10.1109/CDC.2014.7039737
- Roveda, L., Maroni, M., Mazzuchelli, L., Praolini, L., Shahid, A., Bucca, G., et al. (2022). Robot end-effector mounted camera pose optimization in object detection-based tasks. *J. Intelligent Robotic Syst.* 104, 16. doi:10.1007/s10846-021-01558-0
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. (2016). Taking the human out of the loop: a review of bayesian optimization. *Proc. IEEE* 104, 148–175. doi:10.1109/JPROC.2015.2494218
- Siciliano, B., Sciavicco, L., Villani, L., and Oriolo, G. (2008). *Robotics: modelling, planning and control*. 1st edn. Incorporated: Springer Publishing Company.
- Turchetta, M., Berkenkamp, F., and Krause, A. (2019). Safe exploration for interactive machine learning. *Adv. Neural Inf. Process. Syst.* 32.
- Virga, S., Zettinig, O., Esposito, M., Pfister, K., Frisch, B., Neff, T., et al. (2016). “Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms,” in 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Daejeon, Korea, 09–14 October 2016 (IEEE), 508–513. doi:10.1109/IROS.2016.7759101
- Welleweerd, M., de Groot, A., de Looijer, S., Siepel, F., and Stramigioli, S. (2020). “Automated robotic breast ultrasound acquisition using ultrasound feedback,” in 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May 2020 - 31 August 2020 (IEEE), 9946–9952. doi:10.1109/ICRA40945.2020.9196736
- Yi, Z., Calandra, R., Veiga, F., van Hoof, H., Hermans, T., Zhang, Y., et al. (2016). “Active tactile object exploration with Gaussian processes,” in 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Daejeon, Korea, 09–14 October 2016 (IEEE), 4925–4930.
- Zhang, F., Bazarevsky, V., Vakunov, A., Tkachenka, A., Sung, G., Chang, C.-L., et al. (2020). MediaPipe hands: on-device real-time hand tracking. arXiv. doi:10.48550/arXiv.2006.10214