

# Closing the Perception-Action Loop for Semantically Safe Navigation in Semi-Static Environments

Jingxing Qian<sup>1</sup>, Siqi Zhou<sup>2</sup>, Nicholas Jianrui Ren<sup>1</sup>, Veronica Chatrath<sup>2</sup>, Angela P. Schoellig<sup>1,2</sup>

**Abstract**—Autonomous robots navigating in changing environments demand adaptive navigation strategies for safe long-term operation. While many modern control paradigms offer theoretical guarantees, they often assume known extrinsic safety constraints, overlooking challenges when deployed in real-world environments where objects can appear, disappear, and shift over time. In this paper, we present a closed-loop perception-action pipeline that bridges this gap. Our system encodes an online-constructed dense map, along with object-level semantic and consistency estimates into a control barrier function (CBF) to regulate safe regions in the scene. A model predictive controller (MPC) leverages the CBF-based safety constraints to adapt its navigation behaviour, which is particularly crucial when potential scene changes occur. We test the system in simulations and real-world experiments to demonstrate the impact of semantic information and scene change handling on robot behavior, validating the practicality of our approach.

## I. INTRODUCTION

Autonomous robots face increasing demands for long-term navigation in challenging environments like warehouses, office spaces and public roads, where prior knowledge is limited. Accurate mapping and localization are critical for safe and efficient operation in such scenarios. However, most localization and mapping systems assume a static world, which is rarely the case in reality [1]. In these environments, robots may encounter both highly dynamic objects, like humans and other robots, and semi-static objects, such as furniture and pallets, which can appear, disappear and shift over time. Semi-static changes are especially hard to detect, and their mishandling can lead to corrupted maps and lost localization, resulting in catastrophic failure such as obstacle collisions. The growing demand for long-term autonomy necessitates robust and adaptive navigation strategies.

Recent efforts have addressed object-aware localization and mapping in semi-static scenes [2], [3], [4], [5]. These approaches all estimate a 0-to-1 consistency score for each mapped object based on sensor data, updating the changed regions and objects when necessary. However, these works solely focus on mapping and localization accuracy, leaving a notable gap in leveraging the resulting object-level semantic and geometric information for downstream decision-making tasks to ensure safe navigation.

<sup>1</sup>The authors are with the University of Toronto Institute for Aerospace Studies and the University of Toronto Robotics Institute.  
Email: {firstname.lastname}@robotics.utoronto.ca

<sup>2</sup>The authors are with the Technical University of Munich and the Munich Institute of Robotics and Machine Intelligence (MIRMI).  
Emails: {firstname, lastname}@tum.de

This work was supported by the Vector Institute for Artificial Intelligence in Toronto and the NSERC Canadian Robotics Network (NCRN).

A video description of our framework with additional experiments is available at <http://tiny.cc/obj-mpc-cbf>.

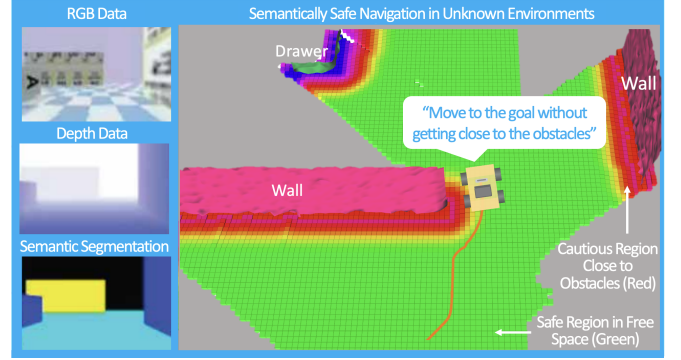


Fig. 1. Our system takes in semantically annotated RGB-D frames to localize and construct an object-aware volumetric map. Each object holds a semantic label and a consistency score. The map is transformed into a CBF, which is utilized by an MPC to plan safe paths around objects. In our framework, objects that are *likely-static*, such as walls, and those having higher consistency scores, will have smaller unsafe regions around them.

An increasing amount of work has been proposed to address the challenge of ensuring safe decision-making for robots operating in changing and uncertain environments [6]. Examples include learning-based model predictive control (MPC) [7], [8], learning-enhanced adaptive control [9], [10], as well as control barrier function (CBF)-based safety filter designs [11], [12]. While these control approaches offer promising theoretical guarantees, a notable limitation lies in their common assumption that safety constraints are known ahead of time [6]. Many practical robot systems rely on on-board sensors to perceive their surroundings, while extrinsic constraints are inferred in real-time. Though there exist a few perception-based safe learning frameworks (e.g., [13], [14]), they do not yet exploit object-level semantic and geometric understanding of the environment that can be readily distilled from perception and mapping systems. In this work, we aim to bridge this gap and facilitate the practical application of theoretical advancements in safe control within real-world contexts. Our emphasis lies in leveraging object-level semantic and geometric constraints to enhance the adaptability and safety of robotic systems operating in semi-static environments.

We present a closed-loop pipeline that establishes an integrated perception-action loop. Our system adopts the state-of-the-art localization and semi-static mapping strategies [3], [15] to construct an up-to-date, object-aware volumetric map (using truncated signed distance functions, TSDFs) of the environment on the fly. Each object within the map maintains a semantic label and a consistency estimate. Upon each update cycle, the map is distilled into a CBF, which synthesizes both the object semantic labels and consistency estimates to regulate safe regions within the scene. Our

system further incorporates an MPC that utilizes the CBF-based safety constraints to adapt its navigation strategy around objects. This adaptability is particularly critical when potential scene changes are detected. We conduct tests in both simulated and real-world environments. Through the experiments, we illustrate how using semantic information and accommodating potential scene changes significantly influence the robot’s behavior, showcasing the practicality of our approach. Our key contributions are as follows:

- We develop a method to encode dense volumetric maps, object semantic labels, and object consistency estimates into a single CBF.
- We build a closed-loop system combining localization, object-aware mapping, and CBF-based safe control for adaptive robot navigation in semi-static scenes.
- We demonstrate the system’s adaptability in both simulated and real-world environments, showcasing dynamic behavior adjustments based on prior semantic knowledge and potential scene changes.

## II. RELATED WORKS

### A. Scene Representations and Object-Aware Mapping

Many geometric representations exist for scene reconstruction, including sparse features from vision systems [15], surfels [16], triangular meshes [17], point clouds [18], and dense volumetric methods such as occupancy grids [19], and Signed Distance Functions (SDFs) [20]. SDFs encode the distance to the nearest surface, gaining traction as they contain rich geometric information for downstream decision-making tasks. SDFs also exploit parallelism to achieve efficient, real-time map updates [20]. Nevertheless, these methods focus on capturing precise scene geometry, ignoring semantic and object-level knowledge.

Researchers have attempted to fuse semantic and object-level information into the geometric map. SLAM++ [21] is an early method which introduces objects into localization and mapping by leveraging existing CAD models from a predefined object library. Voxblox++ [22], Kimera [23], and Hydra [24] build on [20] by fusing semantic labels into the scene-level grid representation. Works like Fusion++ [25], SemanticFusion [26], and MaskFusion [27] leverage semantic segmentation to decompose the scene into object-centric local maps, achieving precise object reconstruction. However, many approaches still assume a static world, rendering them ineffective during real-world deployment where objects may change over time.

### B. Change Detection in Mapping and Localization

To handle dynamics in the scene, two common strategies have been adopted in prior works. The first approach relies on semantic segmentation to identify potentially moving objects, treating them as outliers during camera tracking and scene reconstruction [28], [29], [30], [31], [32]. Although this approach is highly efficient, it can fail when many potentially changing objects are present. The second approach jointly tracks the camera and potentially dynamic objects [27], [33], [34], [35]. These methods rely on motion consistency between consecutive frames to detect high-rate changes, thus not effective when long-term changes are encountered.

Detecting semi-static changes is crucial for safe long-term autonomy, but has been overlooked in literature. Existing methods estimate a geometric consistency score for each feature or object in the scene. Panoptic Multi-TSDF [2] and VI-MID [36] calculate the object consistency based on the overlap ratio between measurements and mapped objects. NeuSE [4], NDFs [37], and 3D-VSG [38] build local scene graphs to predict object-level changes. POCD [3] and Rosen *et al.* [39] propose Bayesian filters to estimate the object stationarity. POV-SLAM [5] proposes an expectation-maximization method to jointly estimate the robot state and object consistency, but incurs high computation costs.

### C. Safe Decision-Making in Robotics

Control theory is a basis for providing safety guarantees based on knowledge of the robot system [6]. There has been an increasing trend of combining control theoretic frameworks and learning-based methods for safe decision-making under uncertainties (e.g., Gaussian Process (GP)-MPC [8], GP-based robust control [40], neural adaptive control [41], and learning-based CBF safety filter [42]). The core idea is enforcing convergence positive invariance of a set that can be the equilibrium in stabilization tasks [43], or a state constraint set in more general settings [8], [42]. Whether using classical or more advanced control techniques for safe decision-making, theoretical guarantees are often provided for constraint sets that are designed a priori [6].

In practice, robots must perceive the environment based on high-dimensional sensory data and compute safe actions from perception. There are a few safe decision-making works that consider perception inputs. For instance, [44] proposes a measurement-robust CBF formulation to account for perception errors, [13] leverages differentiable optimization tools to learn a CBF safety filter for collision avoidance in an end-to-end pipeline, [14] proposes a pipeline for safe locomotion based on plane segmentation, and [45] introduces a dynamic CBF for avoiding dynamic obstacles based on LiDAR measurements. The perception-based approaches, however, have not fully exploited object-level semantic and geometric information that can be obtained from a perception system. In this work, we aim to close the perception and action loop for safe navigation in semi-static environments.

## III. APPROACH OVERVIEW

We explore the closed-loop perception and control problem in semi-static indoor environments, where the robot must simultaneously localize itself, maintain an up-to-date object-level map, track potential changes in the scene, and plan a safe trajectory to reach a desired target pose. This section presents an overview of our proposed approach. A high-level flow diagram of the approach is shown in Figure 2.

### A. Robot System

We consider a ground robot that can be modelled as follows:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t),$$

where  $t \in \mathbb{Z}_{\geq 0}$  is the discrete-time index,  $\mathbf{x} = [x, y, \theta]^T \in \mathcal{X} \subset \mathbb{R}^3$  is the 3-DoF global pose (position and orientation) of the robot with  $\mathcal{X}$  denoting the set of admissible states,

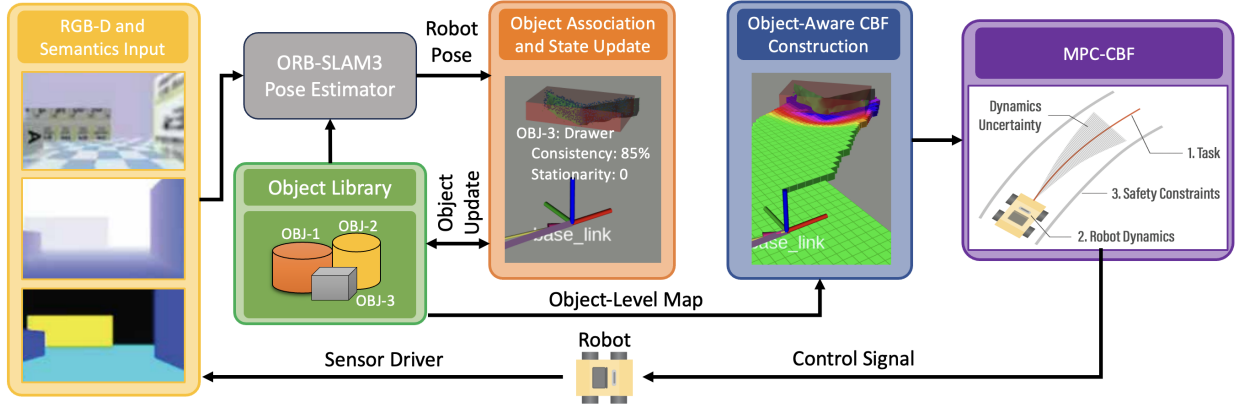


Fig. 2. **Our closed-loop pipeline** (Section III) for semantically safe navigation in semi-static environments. The system maintains an object library and takes in semantically segmented RGB-D frames at each timestep. A modified ORB-SLAM3 provides pose estimates using features from *likely-static* objects. From the estimated pose, observations are associated to mapped objects and the consistency is updated for each object (Section IV-C). From the joint object map, an object-aware CBF is constructed, that takes object semantics and consistency into account (Section IV-E). Finally, an MPC leverages the CBF-based safety constraints to compute actions for the robot (Section IV-F).

$\mathbf{u} = [v_x, v_y, \omega]^T \in \mathcal{U} \subset \mathbb{R}^3$  is the desired linear and angular velocity in the world frame to the robot low-level controller with  $\mathcal{U}$  denoting the set of admissible inputs,  $f: \mathcal{X} \times \mathcal{U} \mapsto \mathcal{X}$  is a Lipschitz continuous function. In this work, we assume that the sets  $\mathcal{X}$  and  $\mathcal{U}$  are convex polytopic sets.

### B. Perception and Mapping

The robot is expected to navigate around rigid objects in a semi-static scene. It is equipped with an RGB-D camera; at each timestep it outputs a frame,  $\mathbf{F}_t$ , that contains an RGB image and an aligned depth image. We further segment  $\mathbf{F}_t$  into semantic pointcloud observations  $\mathcal{Y}_t = \{\mathbf{Y}_{t,j}\}_{j=1\dots J}$ . Given  $\mathcal{Y}_t$ , the localization and mapping systems estimate the current robot pose,  $\hat{\mathbf{x}}_t$ , and update a mapped object library,  $\mathcal{O} = \{\mathbf{O}_i\}_{i=1\dots I}$ , as discussed in Section IV-A-IV-C.

We assume high-level knowledge of the objects is available. Each object belongs to a semantic class,  $c \in \{1 \dots K\}$ , with an associated stationarity class (likelihood of change),  $s \in \{0, 1\}$ , where  $s = 1$  denotes a *likely-static* object (e.g., wall), and  $s = 0$  denotes a *likely-dynamic* object (e.g., chair).

### C. Safe Navigation

Given the tracked objects,  $\mathcal{O}$ , safety constraints are constructed based on the semantic and geometric information to ensure (i) static obstacles are avoided and (ii) conservative behaviours are achieved around inconsistent or *likely-dynamic* objects. In our work, we encode the safety constraint as a CBF. Based on the CBF, the controller computes actions to minimize the distance to the target while adhering to the map-based safety constraints.

## IV. METHODOLOGY

### A. Scene Representation

As our system operates in evolving environments where objects can change over time, it is essential to construct the dense map at an object-level. Following POCD [3], a recent work on mapping in semi-static scenes, each object,  $\mathbf{O}_i$ , in the object library,  $\mathcal{O}$ , consists of the following elements:

- a global 3D position,  $\mathbf{p}_i$ , and 1D heading,  $\phi_i$ ,
- an inferred semantic class,  $c_i$ , and stationarity class,  $s_i$

- a state distribution,  $p(l_i, v_i)$ , to model the object's current geometric change,  $l_i \in \mathbb{R}$ , and consistency (likelihood that the object has not changed),  $v_i \in [0, 1]$ ,
- a TSDF reconstruction,  $\mathbf{M}_i$ , in the global frame.

In our indoor test scenarios, objects are restricted to rotate around the  $z$ -axis. The object TSDF encodes the approximated distance in the global frame to the object surface, with positive values for points outside the object and negative values for points inside the object. All objects in  $\mathcal{O}$  jointly represent the scene, and a trilinearly interpolated 3D TSDF map,  $\mathcal{M}_{tsdf}: \mathbb{R}^3 \mapsto \mathbb{R}$ , can be obtained by overlapping object-level TSDFs and taking the minimum at each voxel.

### B. Localization and Object-Level Mapping

We build our system on top of ORB-SLAM3 [15], a state-of-the-art feature-based SLAM pipeline, and POCD [3]. The system starts with an empty object library,  $\mathcal{O} = \emptyset$ . At each timestamp, the RGB-D frame,  $\mathbf{F}_t$ , and 3D observations,  $\mathcal{Y}_t$ , are input to our modified ORB-SLAM3 to estimate the current robot pose,  $\hat{\mathbf{x}}_t$ . To account for potential scene changes, only feature points belonging to *likely-static* observations (e.g., walls) are used in the estimation.

With the estimated robot pose,  $\hat{\mathbf{x}}_t$ , the observations,  $\mathcal{Y}_t$ , are associated with the mapped objects in  $\mathcal{O}$  using the Hungarian algorithm based on geometric and semantic consistency. Unmatched observations are used to spawn new objects. Each matched observation is integrated into the associated object, and the object's state model,  $p(l, v)$ , is propagated. Note that we employ a decoupled localization and mapping approach due to its high efficiency for real-time performance. A brief overview of the state propagation is provided in Section IV-C, and more details of the mapping pipeline are available in [3].

### C. Object Consistency Estimation in Semi-Static Scenes

We follow POCD [3] to propagate the object state model and estimate the consistency of potentially-changing objects, as it has shown to provide more consistent and noise-robust 3D mapping when faced with discrete scene changes. Specifically, each mapped object,  $\mathbf{O}$ , maintains a joint distribution,

$p(l, v)$ , where  $v \in (0, 1)$  models the likelihood of the object being consistent with historical observations, and  $l \in \mathbb{R}$  models the magnitude of geometric change. In our implementation  $l$  estimates the average TSDF change between the initial reconstruction and subsequent measurements.

When a new observation is available, a geometric measurement,  $p(\Delta | l, v)$ , and a semantic measurement,  $p(s | v)$ , are constructed. Here,  $\Delta \in \mathbb{R}$  is the average TSDF difference between the object reconstruction and its current depth observation, and  $s$  encodes our knowledge on the object stationarity (e.g.,  $s = 1$  if the object is *likely-static*, such as a table, and  $s = 0$  if the object is *likely-dynamic*, such as a cart). These two measurements are used to update the object state model via a Bayesian update rule derived in [3].

In this work, we use both the consistency,  $\mathbb{E}[v]$ , and stationarity label,  $s$ , to regulate the CBF around objects.

#### D. Control Barrier Function (CBF)

CBF certification [11] provides a means to guarantee the safety of a system in the sense of positive invariance of a state constraint set. Let  $h$  be a continuously-differentiable function. Consider a constraint set,  $\mathcal{C}$ , parameterized as the zero-superlevel set of  $h$ . The constraint set and its boundary,  $\partial\mathcal{C}$ , are defined as  $\mathcal{C} = \{\mathbf{x} \in \mathcal{X} | h(\mathbf{x}) \geq 0\}$  and  $\partial\mathcal{C} = \{\mathbf{x} \in \mathcal{X} | h(\mathbf{x}) = 0\}$ , respectively.

For a continuous-time system, a continuously-differentiable function,  $h : \mathcal{X} \mapsto \mathbb{R}$ , is a CBF of the system if (i)  $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}, \forall \mathbf{x} \in \partial\mathcal{C}$ , and (ii) there exists an extended class- $\mathcal{K}_\infty$  function,  $\gamma : \mathbb{R} \mapsto \mathbb{R}$ , such that the following condition is satisfied [11]:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}) \geq -\gamma(h(\mathbf{x})), \quad \forall \mathbf{x} \in \mathcal{X},$$

where  $\dot{h}(\mathbf{x}, \mathbf{u}) = \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \dot{\mathbf{x}} = \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) f_c(\mathbf{x}, \mathbf{u})$  with  $\dot{\mathbf{x}} = f_c(\mathbf{x}, \mathbf{u})$  being the continuous-time model of the system.

Given a CBF, the following condition is often imposed as a constraint in the safety filter design to guarantee positive invariance of the constraint set:

$$\dot{h}(\mathbf{x}, \mathbf{u}) \geq -\gamma(h(\mathbf{x})).$$

For a discrete-time system, a similar CBF condition can be defined as follows:

$$\Delta h(\mathbf{x}_t, \mathbf{u}_t) \geq -\bar{\gamma}h(\mathbf{x}_t),$$

where  $\Delta h(\mathbf{x}_t, \mathbf{u}_t) = h(f(\mathbf{x}_t, \mathbf{u}_t)) - h(\mathbf{x}_t)$ , and  $\bar{\gamma} \in (0, 1]$  is a constant scalar [46].

#### E. Bridging Object Mapping and CBFs

1) *Non-Semantic CBF Construction*: The object TSDFs used to construct the global map,  $\mathcal{M}_{tsdf}$ , are based on ray-tracing from estimated robot poses, causing bias around object boundaries with limited views. Following prior work [45], [47], [48], we generate a 2.5D non-semantic Euclidean Distance Function (EDF) map, denoted as  $\tilde{\mathcal{M}}_{edf} : \mathbb{R}^2 \mapsto \mathbb{R}$ . To capture the obstacles within the robot's traversal space, we first extract the TSDF layers between the ground plane and a specified height threshold,  $\theta_z$ . We concatenate the layers into a single 2.5D map by taking the minimum of the absolute value along the  $z$ -axis.

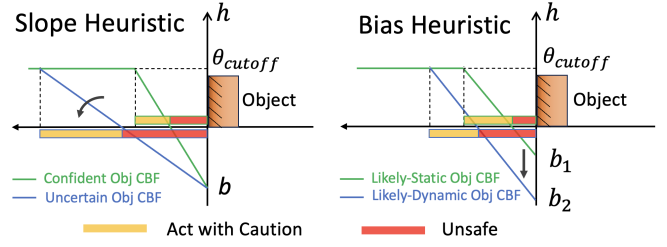


Fig. 3. Visualization of the proposed heuristics for CBFs around objects. The Slope Heuristic scales down the slope of the CBF for uncertain objects, inflating both the unsafe region (red) and the cautious region (orange), causing the robot to act more conservatively. The Bias Heuristic only inflates the unsafe region around *likely-dynamic* objects to create a larger safety buffer, and the robot's behavior is unaffected outside. The two heuristics can be combined to achieve the desired behaviours.

The approximated zero-level set,  $\partial\tilde{\mathcal{M}}_{edf}$ , thresholded by a small value,  $\theta_{zero}$ , representing the obstacle surfaces, is extracted, and the Euclidean distance is computed for  $\tilde{\mathcal{M}}_{edf}$ . The process is formalized as follows:

$$\mathcal{M}_{2.5D}(x, y) = \min_{0 < z \leq \theta_z} |\mathcal{M}_{tsdf}(x, y, z)| \quad (1a)$$

$$\partial\tilde{\mathcal{M}}_{edf} = \{(x, y) | \mathcal{M}_{2.5D}(x, y) \leq \theta_{zero}\} \quad (1b)$$

$$\tilde{\mathcal{M}}_{edf}(x, y) = \min_{(\bar{x}, \bar{y}) \in \partial\tilde{\mathcal{M}}_{edf}} \|(x, y) - (\bar{x}, \bar{y})\|_2 \quad (1c)$$

We apply two transformations to encode  $\tilde{\mathcal{M}}_{edf}$  into a CBF. To account for the size of the robot, we apply a negative bias,  $b$ , which introduces a buffer around the obstacles. To ensure the robot's behaviour remains unaffected when far from obstacles, we truncate  $\tilde{\mathcal{M}}_{edf}$  by a cutoff threshold,  $\theta_{cutoff}$ . Formally, the continuous non-semantic CBF, denoted by  $\tilde{h}$ , is defined as:

$$\tilde{h}(\mathbf{x}) = \min\{\tilde{\mathcal{M}}_{edf}(x, y) - b, \theta_{cutoff}\} \quad (2)$$

In our implementation, the continuity is achieved via bilinear interpolation over a discrete 2D grid, and gradients of  $\tilde{h}$  is obtained via finite difference.

#### 2) Incorporating Semantic and Geometric Knowledge:

The previously defined CBF,  $\tilde{h}$ , does not consider any object-level semantics and consistency knowledge, which may lead to risky behaviour such as moving too close to potentially changed objects. We propose the Slope Heuristic and the Bias Heuristic to generate an object-aware EDF map,  $\mathcal{M}_{edf}$ , for enhanced safety and adaptability. We first augment the zero-level set,  $\partial\tilde{\mathcal{M}}_{edf}$ , with the object consistency,  $\mathbb{E}[v]$ , and stationarity class,  $s$ , for which the voxel belongs to:

$$\partial\mathcal{M}_{edf} = \{(x, y, \mathbb{E}[v_i], s_i) | \mathcal{M}_{2.5D}(x, y) \leq \theta_{zero}, (x, y) \in \mathbf{O}_i\} \quad (3)$$

The Slope Heuristic scales the Euclidean distances in (1c) to regulate both the permissible proximity of the robot to uncertain objects and the aggressiveness when navigating around them. Under the common bias,  $b$ , and cutoff threshold,  $\theta_{cutoff}$ , CBFs with lower slopes will have inflated unsafe regions, and the robot will behave more conservatively when approaching them. We propose to utilize the adjusted object-level consistency estimate,  $\lambda_c \mathbb{E}[v]$ , as the scaling factor, where  $\lambda_c$  is a tunable parameter.



The Bias Heuristic changes the bias to solely control the unsafe region around objects. We propose to upscale  $b$  with a tunable parameter,  $\lambda_s > 1$ , for objects that are *likely-dynamic* ( $s=0$ ). This ensures that the robot does not get too close to potentially changing objects, but can also move freely, when having a high level of confidence. Combining the two heuristics, our object-aware CBF,  $h$ , is formulated as

$$\begin{aligned} \mathcal{M}_{edf}(x, y) &= \min_{(\bar{x}, \bar{y}, \mathbb{E}[v], s) \in \partial \mathcal{M}_{edf}} \lambda_c \mathbb{E}[v] \|(x, y) - (\bar{x}, \bar{y})\|_2 \\ &\quad - [\lambda_s(1-s) + s]b \end{aligned} \quad (4)$$

$$h(\mathbf{x}) = \min\{\mathcal{M}_{edf}(x, y), \theta_{\text{cutoff}}\} \quad (5)$$

We visualize the effect of the two heuristics in Figure 3.

#### F. Discrete-Time Safe Control

In this work, inspired by [46], we use an MPC-CBF framework for computing safe actions based on the object-level map. At each time,  $t$ , given the current state,  $\mathbf{x}_t$ , and the CBF,  $h$ , the following optimization problem is solved:

$$\min_{\substack{\mathbf{x}_{t+1:t+T}|t \\ \mathbf{u}_{t:t+T-1}|t}} \sum_{k=0}^{T-1} l_t(\mathbf{x}_{t+k}|t, \mathbf{u}_{t+k}|t) + l_T(\mathbf{x}_{t+T}|t) \quad (6a)$$

$$\text{s.t. } \mathbf{x}_t|t = \mathbf{x}_t, \quad (6b)$$

$$\mathbf{x}_{t+k+1}|t = f(\mathbf{x}_{t+k}|t, \mathbf{u}_{t+k}|t), \forall k \in \mathcal{T}, \quad (6c)$$

$$\mathbf{x}_{t+k}|t \in \mathcal{X}, \mathbf{u}_{t+k}|t \in \mathcal{U}, \forall k \in \mathcal{T}, \quad (6d)$$

$$\Delta h(\mathbf{x}_{t+k}|t, \mathbf{u}_{t+k}|t) \geq -\bar{\gamma} h(\mathbf{x}_{t+k}|t), \forall k \in \mathcal{T}, \quad (6e)$$

where  $T$  is the prediction horizon length,  $\mathcal{T} = \{0, 1, \dots, T-1\}$  is the set of time indices over the prediction horizon, “ $\cdot|t$ ” in the subscripts denotes consecutive timesteps, and  $l_t$  and  $l_T$  are the stage cost and the terminal cost, respectively. After solving the optimization problem in (6), the first input,  $\mathbf{u}_t|t$ , is applied to the system.

In contrast to a typical MPC, the MPC-CBF framework allows us to (i) conveniently define the desired behaviour of the robot via a CBF directly derived based on map information, and (ii) reduces the prediction horizon required to achieve a desired level of safety behaviour [46].

Note that the MPC-CBF is not a standard convex optimization problem due to the dynamics in (6c) and the CBF constraint in (6e). To allow fast online computation, we approximate (6c) and (6e) by linearizing them about the prediction trajectory from the previous timestep. For each timestep,  $k \in \mathcal{T}$ , we introduce a new set of decision variables,  $\delta \mathbf{x}_{t+k}|t = \mathbf{x}_{t+k}|t - \mathbf{x}_{\text{op}}$  and  $\delta \mathbf{u}_{t+k}|t = \mathbf{u}_{t+k}|t - \mathbf{u}_{\text{op}}$ , where  $\mathbf{x}_{\text{op}} = \mathbf{x}_{t-1+k}|t-1$  and  $\mathbf{u}_{\text{op}} = \mathbf{u}_{t-1+k}|t-1$ . The linearized dynamics and CBF constraints can be written in forms that are affine in the decision variables:

$$\delta \mathbf{x}_{t+k+1}|t = \mathbf{A}(\mathbf{x}_{\text{op}}, \mathbf{u}_{\text{op}}) \delta \mathbf{x}_{t+k}|t + \mathbf{B}(\mathbf{x}_{\text{op}}, \mathbf{u}_{\text{op}}) \delta \mathbf{u}_{t+k}|t, \quad (7)$$

$$\mathbf{C}(\mathbf{x}_{\text{op}}, \mathbf{u}_{\text{op}}) \delta \mathbf{x}_{t+k}|t + \mathbf{D}(\mathbf{x}_{\text{op}}, \mathbf{u}_{\text{op}}) \delta \mathbf{u}_{t+k}|t + c(\mathbf{x}_{\text{op}}, \mathbf{u}_{\text{op}}) \geq 0, \quad (8)$$

where  $c$ ,  $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$ , and  $\mathbf{D}$  are constant terms derived from the linear approximations. In our implementation, (8) requires querying  $h$  and its gradient,  $\frac{\partial h}{\partial \mathbf{x}}$ , at

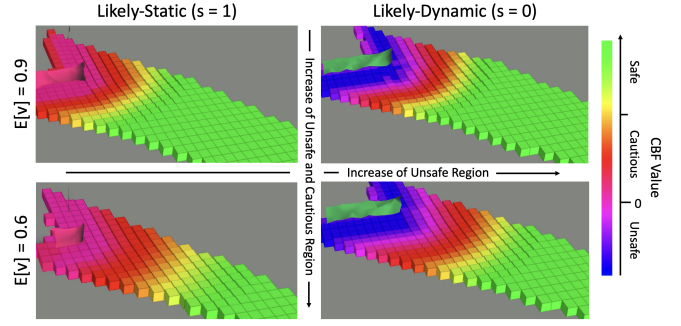


Fig. 4. Visualization of the object-aware CBFs around the *drawer* object under different stationarity labels and consistency estimates.

$\{\mathbf{x}_t^*|t-1, \dots, \mathbf{x}_{t-1+T}^*|t-1\}$ . By replacing the dynamics, (6c), with (7), and the CBF constraint, (6e), with (8), and re-expressing (6a), (6b), and (6d) in terms of  $\delta \mathbf{x}$  and  $\delta \mathbf{u}$ , we obtain a quadratic program that can be solved efficiently.

## V. SIMULATION RESULTS

### A. Experiment Setup

We demonstrate the key features of our object-aware MPC-CBF framework in simulation. As seen in Figure 1, an indoor scenario consisting of multiple *walls* and one *drawer* is set up in PyBullet [49]. The *walls* are considered *likely-static*. An omnidirectional robot, modeled as single integrators in each dimension, is simulated and controlled by our system. The robot is 50 cm in diameter and equipped with a RGB-D sensor that outputs color images, Gaussian corrupted depth measurements, and ground-truth semantic masks. To construct the object-aware CBF, we set the zero-level threshold,  $\theta_{\text{zero}}$ , to 0.15 m, truncation distance,  $\theta_{\text{cutoff}}$ , to 1.8 m, bias,  $b$ , to 0.75 m, consistency factor,  $\lambda_c$ , to 3.0, and stationarity factor,  $\lambda_s$ , to 2.0.

### B. Object Semantic Knowledge and Consistency

We validate our CBF formulation’s capture of object-level semantic knowledge and consistency estimates, as shown in Figure 4. As the robot approaches the stationary *drawer*, its consistency estimate increases with more observations, causing the unsafe (enclosed by pink grids) and the cautious (enclosed by yellow grids) regions to contract. When testing the *drawer* with a *likely-dynamic* label, the Bias Heuristic effectively increases the unsafe region, while the cautious region remains the same (under the same consistency level) due to the unchanged slope of the CBF.

### C. Balancing Feasibility and Conservatism

The scalar,  $\bar{\gamma}$ , in the MPC-CBF safety constraint, (6e), is a critical design variable. A smaller  $\bar{\gamma}$  makes the controller safer but may render the optimization in (6) infeasible. Conversely, a larger  $\bar{\gamma}$  improves controller feasibility but elevates risk. Figure 5 shows robot trajectories under different  $\bar{\gamma}$  choices, with lower values resulting in more conservative behaviour around the wall. Note that with  $\bar{\gamma} = 0.1$ , the robot went straight through the red cautious region but did not hit the pink zero-level set. In real-world experiments, we opt for  $\bar{\gamma} = 0.03$  so the system is not overly conservative under measurement noise and localization errors.

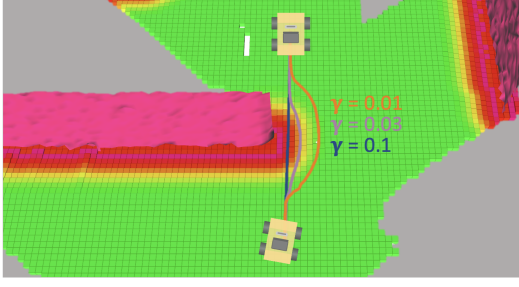


Fig. 5. Ground-truth robot trajectories with different values of  $\bar{\gamma}$  in the MPC-CBF safety constraint (6e).

## VI. REAL WORLD EXPERIMENTS

### A. Experiment Setup

We demonstrate the closed-loop capability of our system on a real-world robot. As in the simulation tests, we use an omnidirectional mobile base equipped with an Intel Realsense D435 RGB-D camera for perception. The pipeline runs on an external laptop with an Intel i7-8850H CPU, at a fixed rate of 5 Hz. A low-level onboard PID controller follows the MPC trajectory at 20 Hz. We set up the test scenario as seen in Figure 6 and adopted the same parameter settings from Section V-A. While we tested on an omnidirectional robot, our method can be generalized to other kinematic models.

### B. Obstacle Avoidance with Semantic Knowledge

In this experiment, the robot navigates through a narrow gap between a wheeled drawer (*likely-dynamic*, left), and a flat-base drawer (*likely-static*, right). We compare our object-aware MPC-CBF against two baselines derived from (6): a naive MPC-CBF using the non-semantic CBF  $\tilde{h}$  in (2), and a classic MPC replacing the full CBF constraint (6e) with geometric state constraints,  $\tilde{h}(\mathbf{x}_{t+k|t}) > 0$ .

Figure 6 depicts the ground-truth trajectories recorded by a Vicon motion capture system. The classic MPC solely considers the state constraint at the unsafe region's boundary. Due to depth measurement noise and localization errors, the unsafe region frequently expands and contracts. As a result, the robot under the MPC often enters infeasible regions, getting stuck between the two drawers.

In contrast, the non-semantic MPC-CBF successfully navigates the gap. Since the controller does not encode any semantic information, the robot opts for the middle path, disregarding potential risks from scene changes.

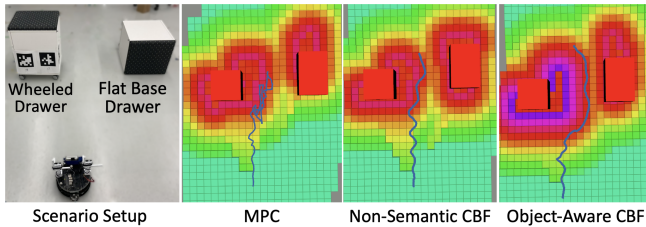


Fig. 6. The scenario (left) and visualization (right) of the ground-truth robot trajectories using three controllers. The wheeled drawer is considered to be likely to change ( $s = 1$ ) while the flat base drawer is likely to stay ( $s = 0$ ). The robot distinguishes between the two drawers based on color.

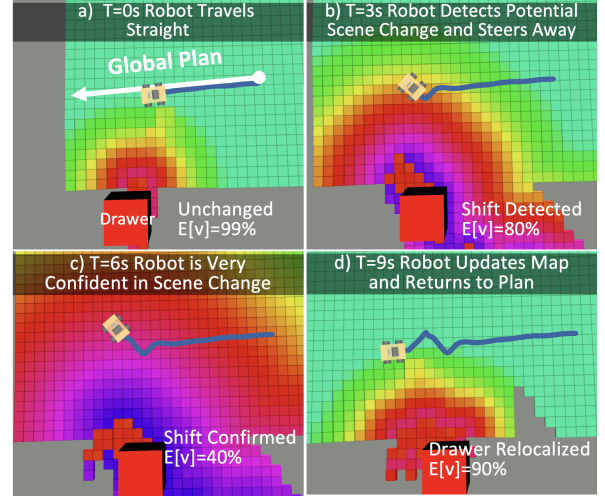


Fig. 7. An illustration of how our closed-loop system deviates from the original path in response to a sudden change in the scene.

Our object-aware MPC-CBF recognizes that the left drawer may undergo changes, resulting in larger unsafe and cautious regions around it. Consequently, the robot navigates the gap while remaining closer to the *likely-static* drawer. Although the scenario is simple, our local planning method can be scaled to more complex settings when combined with a global planner such as [50].

### C. Reacting to Scene Changes

In the final experiment, we showcase the system's adaptability to sudden scene changes, as depicted in Figure 7. Initially, the robot follows a straight path near a drawer (a). When the drawer is pushed, causing a scene change, our semi-static mapping system quickly detects it. The drawer's consistency estimate begins to drop, leading to an expansion of the unsafe region around it (b). The object-aware MPC-CBF responds to this by redirecting the robot away from its initial path (c). Once the drawer's consistency estimate falls below 40%, it is removed from the object library and reconstructed at its new pose with high consistency. As a result, the unsafe region returns to normal, and the robot resumes its initial path, continuing towards its goal.

## VII. CONCLUSION

In this work, we present a system designed for closed-loop perception and safe control within semi-static environments. Our approach involves encoding object-level semantic prior knowledge and consistency estimates into the CBF safety constraints within the MPC-CBF framework. This integration of scene semantics and consistency offers a versatile means to directly shape the robot's behavior in complex environments based on map-derived information. Through a series of experiments conducted in both simulated and real-world scenarios, we demonstrate how our system dynamically adapts its actions to address potential risks in the environment, thereby ensuring safe navigation. By addressing the challenges posed by changing environments, our work presents promising prospects for the development of robust systems capable of navigating extended periods in dynamic real-world settings.

## REFERENCES

- [1] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. D. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, pp. 1309–1332, 2016. [Online]. Available: <https://api.semanticscholar.org/CorpusID:9686224>
- [2] L. M. Schmid, J. A. Delmerico, J. L. Schönberger, J. I. Nieto, M. Pollefeys, R. Siegwart, and C. Cadena, "Panoptic multi-TSDFs: a flexible representation for online multi-resolution volumetric mapping and long-term dynamic scene consistency," *ICRA*, 2022.
- [3] J. Qian, V. Chatrath, J. Yang, J. Servos, A. Schoellig, and S. L. Waslander, "POCD: Probabilistic Object-Level Change Detection and Volumetric Mapping in Semi-Static Scenes," in *2022 Robotics: Science and Systems (RSS)*, 2022.
- [4] J. Fu, Y. Du, K. Singh, J. B. Tenenbaum, and J. J. Leonard, "Neuse: Neural se(3)-equivariant embedding for consistent spatial understanding with objects," in *Robotics: Science and Systems (RSS)*, 2023.
- [5] J. Qian, V. Chatrath, J. Servos, A. Mavrinac, W. Burgard, S. L. Waslander, and A. Schoellig, "POV-SLAM: Probabilistic Object-Level Variational SLAM," in *2023 Robotics: Science and Systems (RSS)*, 2023.
- [6] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.
- [7] U. Rosolia and F. Borrelli, "Learning model predictive control for iterative tasks: a data-driven control framework," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1883–1896, 2017.
- [8] C. J. Ostafew, A. P. Schoellig, T. D. Barfoot, and J. Collier, "Learning-based nonlinear model predictive control to improve vision-based mobile robot path tracking," *Journal of Field Robotics*, vol. 33, no. 1, pp. 133–152, 2016.
- [9] A. Gahlawat, P. Zhao, A. Patterson, N. Hovakimyan, and E. Theodorou, "L1-GP: L1 adaptive control with bayesian learning," in *Proc. of the Learning for Dynamics and Control Conference*, 2020, pp. 826–837.
- [10] G. Chowdhary, H. A. Kingravi, J. P. How, and P. A. Vela, "Bayesian nonparametric adaptive control using Gaussian processes," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 3, pp. 537–550, 2014.
- [11] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. of the European control conference (ECC)*, 2019, pp. 3420–3431.
- [12] L. Wang, E. A. Theodorou, and M. Egerstedt, "Safe learning of quadrotor dynamics using barrier certificates," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2460–2465.
- [13] W. Xiao, T.-H. Wang, R. Hasani, M. Chahine, A. Amini, X. Li, and D. Rus, "Barriernet: Differentiable control barrier functions for learning of safe robot control," *IEEE Transactions on Robotics*, 2023.
- [14] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, "Perceptive locomotion through nonlinear model-predictive control," *IEEE Transactions on Robotics*, 2023.
- [15] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [16] H. Pfister, M. Zwicker, J. van Baar, and M. H. Gross, "Surfels: surface elements as rendering primitives," *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000. [Online]. Available: <https://api.semanticscholar.org/CorpusID:1756681>
- [17] M. Bloesch, T. Laidlow, R. Clark, S. Leutenegger, and A. J. Davison, "Learning meshes for dense visual slam," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5854–5863, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:203600671>
- [18] C. Kerl, J. Sturm, and D. Cremers, "Dense visual slam for rgb-d cameras," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 2100–2106.
- [19] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013.
- [20] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [21] R. F. Salas-Moreno, R. A. Newcombe, H. M. Strasdat, P. H. J. Kelly, and A. J. Davison, "SLAM++: Simultaneous localisation and mapping at the level of objects," *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [22] M. Grinvald, F. Furrer, T. Novkovic, J. J. Chung, C. Cadena, R. Siegwart, and J. Nieto, "Volumetric Instance-Aware Semantic Mapping and 3D Object Discovery," *IEEE Robotics and Automation Letters*, 2019.
- [23] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [24] N. Hughes, Y. Chang, and L. Carlone, "Hydra: A real-time spatial perception system for 3D scene graph construction and optimization," in *2022 Robotics: Science and Systems (RSS)*, 2022.
- [25] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, "Fusion++: Volumetric object-level SLAM," *2018 International Conference on 3D Vision (3DV)*, 2018.
- [26] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3d semantic mapping with convolutional neural networks," *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [27] M. Rünz and L. de Agapito, "Maskfusion: Real-time recognition, tracking and reconstruction of multiple moving objects," *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2018.
- [28] T. Sun, Y. Sun, M. Liu, and D.-Y. Yeung, "Movable-object-aware visual SLAM via weakly supervised semantic segmentation," *ArXiv*, 2019.
- [29] A. Rosinol, M. Abate, J. Shi, and L. Carlone, "3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans," in *Robotics: Science and Systems*, 2020.
- [30] X. Lu, H. Wang, S. Tang, H. Huang, and C. Li, "Dm-slam: Monocular SLAM in dynamic environments," *Applied Sciences*, 2020.
- [31] I. Ballester, A. Fontán, J. Civera, K. H. Strobl, and R. Triebel, "DOT: dynamic object tracking for visual SLAM," *CoRR*, 2020.
- [32] C. Yu, Z. Liu, X. Liu, F. Xie, Y. Yang, Q. Wei, and F. Qiao, "DS-SLAM: A semantic visual SLAM towards dynamic environments," *CoRR*, 2018.
- [33] R. Hachiuma, C. Pirchheim, D. Schmalstieg, and H. Saito, "Detect-fusion: Detecting and segmenting both known and unknown dynamic objects in real-time SLAM," in *BMVC*, 2019.
- [34] M. Grinvald, F. Tombari, R. Y. Siegwart, and J. I. Nieto, "TSDF++: A multi-object formulation for dynamic object tracking and reconstruction," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [35] B. Xu, W. Li, D. Tzoumanikas, M. Bloesch, A. J. Davison, and S. Leutenegger, "MID-fusion: Octree-based object-level multi-instance dynamic SLAM," *2019 International Conference on Robotics and Automation (ICRA)*, 2019.
- [36] Y. Ren, B. Xu, C. L. Choi, and S. Leutenegger, "Visual-inertial multi-instance dynamic SLAM with object-level relocation," in *arXiv*, 2022.
- [37] J. Fu, Y. Du, K. Singh, J. B. Tenenbaum, and J. J. Leonard, "Robust change detection based on neural descriptor fields," *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2817–2824, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:251223475>
- [38] S. Looper, J. R. Puigvert, R. Y. Siegwart, C. Cadena, and L. M. Schmid, "3d vsg: Long-term semantic scene change prediction through 3d variable scene graphs," *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8179–8186, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252355114>
- [39] D. M. Rosen, J. Mason, and J. J. Leonard, "Towards lifelong feature-based mapping in semi-static environments," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016.
- [40] F. Berkenkamp and A. P. Schoellig, "Safe and robust learning control with gaussian processes," in *Proc. of the European Control Conference (ECC)*, 2015, pp. 2496–2501.
- [41] G. Joshi and G. Chowdhary, "Deep model reference adaptive control," in *Proc. of the IEEE Conference on Decision and Control (CDC)*, 2019, pp. 4601–4608.

- [42] A. Taylor, A. Singletary, Y. Yue, and A. Ames, "Learning for safety-critical control with control barrier functions," in *Proc. of the Learning for Dynamics and Control Conference*, 2020, pp. 708–717.
- [43] S. M. Richards, F. Berkenkamp, and A. Krause, "The Lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems," in *Proc. of the Conference on Robot Learning*, 2018, pp. 466–476.
- [44] S. Dean, A. Taylor, R. Cosner, B. Recht, and A. Ames, "Guaranteeing safety of learned perception modules via measurement-robust control barrier functions," in *Proc. of the Conference on Robot Learning*, 2021, pp. 654–670.
- [45] Z. Jian, Z. Yan, X. Lei, Z.-R. Lu, B. Lan, X. Wang, and B. Liang, "Dynamic control barrier function-based model predictive control to safety-critical obstacle-avoidance of mobile robot," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3679–3685, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252367490>
- [46] J. Zeng, B. Zhang, and K. Sreenath, "Safety-critical model predictive control with discrete-time control barrier function," in *Proc. of the American Control Conference (ACC)*, 2021, pp. 3882–3889.
- [47] A. W. Singletary, K. Klingebiel, J. R. Bourne, A. W. Browning, P. T. Tokumaru, and A. D. Ames, "Comparative analysis of control barrier functions and artificial potential fields for obstacle avoidance," *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8129–8136, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:224803757>
- [48] L. Han, F. Gao, B. Zhou, and S. Shen, "Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots," *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4423–4430, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:70350033>
- [49] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [50] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. J. Teller, "Anytime motion planning using the rrt\*," *2011 IEEE International Conference on Robotics and Automation*, pp. 1478–1483, 2011. [Online]. Available: <https://api.semanticscholar.org/CorpusID:1003020>