

Learning Barrier Functions with Memory for Robust Safe Navigation

Kehan Long* Cheng Qian* Jorge Cortés Nikolay Atanasov

Abstract—Control barrier functions are widely used to enforce safety properties in robot motion planning and control. However, the problem of constructing barrier functions online and synthesizing safe controllers that can deal with the associated uncertainty has received little attention. This paper investigates safe navigation in unknown environments, using onboard range sensing to construct control barrier functions online. To represent different objects in the environment, we use the distance measurements to train neural network approximations of the signed distance functions incrementally with replay memory. This allows us to formulate a novel robust control barrier safety constraint which takes into account the error in the estimated distance fields and its gradient. Our formulation leads to a second-order cone program, enabling safe and stable control synthesis in a priori unknown environments.

I. INTRODUCTION

Modern applications of ground, aerial, and underwater mobile robots necessitate safe operation in unexplored and unstructured environments. Planning and control techniques that ensure safety in such environments, relying on limited field-of-view observations, are critical for autonomous navigation. A successful and widely used approach for safe control of dynamical systems is to encode safety constraints through control barrier functions (CBFs) [1] and combine the CBF constraints with the control objectives in a quadratic programming (QP) formulation. A common assumption in specifying CBF safety constraints is that the robot already has complete knowledge of the unsafe regions. However, navigation in unknown environments requires the online estimations of unsafe regions using onboard sensing.

This paper studies a navigation problem in which a robot, equipped with a LiDAR scanner, is tasked to follow a motion plan relying on streaming range measurements to ensure safety. A control Lyapunov function (CLF) [1] is utilized to specify the trajectory tracking objective. We construct CBFs approximating the shape of the objects in the environment and introduce associated robust safety constraints in the optimization of the robot’s control policy to ensure obstacle avoidance.

The signed distance function (SDF) of a set Ω in a metric space determines the distance of a given point x to the boundary of Ω , with the sign determined by whether or not x is in Ω . Observing that CBF and SDF are both level-set functions, we approximate the SDFs of observed objects and define the corresponding CBFs. Using the LiDAR measurements at each time instant, we incrementally train a multi-

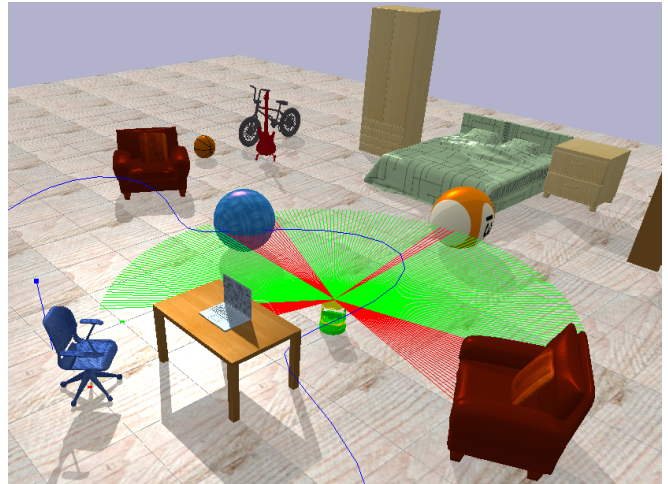


Fig. 1: A ground robot equipped with a LiDAR scanner is navigating safely along a desired path (blue curve) in an unknown room.

layer perceptron (MLP) to model each obstacle’s SDF shape in real time. To balance the accuracy and efficiency of SDF reconstruction, we employ replay memory in training [2].

After obtaining SDF approximations of the obstacles, we can efficiently estimate the distance and gradient information for the barrier functions. We take into account the errors in the estimation for the controller synthesis, resulting in safety constraints where the input appears linearly and also within an l_2 -norm term. The presence of such constraints means that the control synthesis problem can no longer be stated as a QP. Nevertheless, we show that the resulting optimization problem can be formulated as a second-order cone programming (SOCP) problem and therefore can still be solved efficiently. The proposed CLF-CBF-SOCP framework integrates the SDF estimation of the CBFs to synthesize safe controls at each time instant, allowing the robot to safely navigate the unknown environment. This paper makes the following contributions.

- An online incremental training approach, utilizing replay memory, is proposed for deep neural network approximation of the signed distance functions of objects observed with streaming distance measurements.
- Our analysis shows that incorporating safety constraints, that account for the worst-case estimation error of the SDF approximations, into a control synthesis optimization problem leads to a convex SOCP formulation.

II. RELATED WORK

Occupancy mapping techniques aim to partition an environment into occupied and free subsets using sensor observations. Octomap [3] is a probabilistic occupancy mapping

* indicates equal contribution.

We gratefully acknowledge support from NSF RI IIS-2007141.

The authors are with the Contextual Robotics Institute, University of California San Diego, La Jolla, CA 92093, USA. {k3long, chqian, cortes, natanasov}@ucsd.edu

algorithm that uses a tree-based structure to store and update the probability of occupied and free space based on streaming range observations. Occupancy information alone may not be enough for safe planning and control, since many algorithms require distance (or even gradient) information to the obstacle surfaces. Voxblox [4] is an incremental truncated SDF mapping algorithm that approximates the distance to the nearest surface at a set of weighted support points. The approach also provides an efficient extension from truncated to complete (Euclidean) SDF using an online wavefront propagation algorithm. Fiesta [5] further improves the efficiency and accuracy of building online SDF maps by using two independent queues for inserting and deleting obstacles separately and doubly linked lists for map maintenance. Recently, impressive results have been achieved in object shape modeling using deep neural networks. DeepSDF [6] and Occupancy Nets [7] implicitly represent 3D shapes by supervised learning of SDF approximations via fully connected neural networks. Gropp *et al.* [8] further extended the DeepSDF model by introducing the *Eikonal* constraint that any SDF function must satisfy in training the neural network. In our paper, we further extend this offline training method in two ways: enabling online training by introducing replay memory, which helps the network to remember the reconstructed shape from past observations, and adding truncated signed distance points to the training set to ensure the correct sign of the approximated SDF.

Quadratic programming with CBF constraints offers an elegant and efficient framework for safe control synthesis in robot navigation tasks [9], [10]. This approach employs CLF constraints to stabilize the system and achieving the control objectives whenever possible, whereas the CBF constraints are used to ensure safety of the resulting controller at all times. CBFs are used for safe multi-robot navigation in [9]. CLF and CBF constraints are combined to solve a simultaneous lane-keeping (LK) and adaptive speed regulation (ASR) problem in [10]. Here, CLFs are used to ensure convergence to the control objectives for LK and ASR, where CBFs are used to meet safety requirements. However, the barrier functions in these approaches are assumed to be known. When the environment is unknown, a robot may only rely on the estimation of barrier functions using its sensors. A closely related work by Srinivasan *et al.* [11] presents a Support Vector Machine (SVM) approach for the online synthesis of barrier functions from range data. By approximating the boundary between occupied and free space as the SVM decision boundary, the authors extract CBF constraints and solve a CBF-QP to generate safe control inputs. Our approach constructs CBFs directly from an approximation of the object SDFs and explicitly considers the potential errors in the CBF constraints when formulating the control synthesis optimization problem.

III. PROBLEM FORMULATION

Consider a robot whose dynamics are governed by a nonlinear control-affine system:

$$\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}, \quad (1)$$

where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n$ is the robot state, $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ and $g : \mathbb{R}^n \mapsto \mathbb{R}^{n \times m}$ are continuously differentiable functions, and $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$ is the control input. Define the admissible control input space as $\mathcal{U} := \{\mathbf{u} \in \mathbb{R}^m \mid \mathbf{A}\mathbf{u} \leq \mathbf{b}\}$, where $\mathbf{A} \in \mathbb{R}^{k \times m}$ and $\mathbf{b} \in \mathbb{R}^k$. The state space \mathcal{X} is partitioned into a closed safe set \mathcal{S} and open obstacle set \mathcal{O} such that $\mathcal{S} \cap \mathcal{O} = \emptyset$ and $\mathcal{X} = \mathcal{S} \cup \mathcal{O}$. The obstacle set is further partitioned into K obstacles, denoted $\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_K$, such that $\mathcal{O} = \bigcup_{i=1}^K \mathcal{O}_i$ and $\mathcal{O}_j \cap \mathcal{O}_k = \emptyset$ for any $0 \leq j, k \leq K$. Each of these sets is defined implicitly through a continuously differentiable function $\varphi_i : \mathcal{X} \mapsto \mathbb{R}$, with gradient $\nabla \varphi_i : \mathcal{X} \mapsto \mathbb{R}^n$. Formally, $\mathcal{O}_i = \{\mathbf{x} \in \mathbb{R}^n \mid \varphi_i(\mathbf{x}) < 0\}$. Note that $\varphi_i(\mathbf{x}) < 0$ if the robot state \mathbf{x} is in the open set \mathcal{O}_i , and $\varphi_i(\mathbf{x}) = 0$ if \mathbf{x} is on its boundary $\partial \mathcal{O}_i$.

The robot is equipped with a range sensor, such as a LiDAR scanner, and aims to follow a desired path, relying on the noisy distance measurements to avoid collisions. The path and sensed information are defined as follows.

Definition 1. A *path* is a piece-wise continuous function $r : [0, 1] \mapsto \text{Int}(\mathcal{S})$.

Definition 2. Let $\mathcal{F}(\mathbf{x}) \subset \mathcal{X}$ be the field of view of the range sensor when the robot is in state \mathbf{x} . At discrete times t_k , for $k \in \mathbb{N}$, the range sensor provides a set of points $\mathcal{P}_{k,i} = \{\mathbf{p}_{k,i,j}\}_j \subset \mathcal{F}(\mathbf{x}(t_k)) \cap \partial \mathcal{O}_i$ on the boundary of each obstacle \mathcal{O}_i which are within its sensing range.

Problem. Consider the system in (1) operating in an unknown environment, i.e., the obstacles sets $\{\mathcal{O}_i\}_{i=1}^K$ are unknown a priori. Given a desired path r and noisy range sensor measurements $\mathcal{P}_{k,i}$ for $k \in \mathbb{N}$ and $i = 1, \dots, K$, design a feedback controller $\mathbf{x} \mapsto \mathbf{u}(\mathbf{x})$ for (1) that ensures the robot can safely move along the path r , i.e., $\mathbf{x}(t) \in \mathcal{S}$ for all t and $\mathbf{x}(t) \rightarrow r(1)$ as $t \rightarrow \infty$.

IV. OBSTACLE ESTIMATION VIA ONLINE SIGNED DISTANCE FUNCTION APPROXIMATION

Throughout the paper, we rely on the concept of Signed Distance Function (SDF) to describe each unsafe region $\{\mathcal{O}_i\}_{i=1}^K$. For each i , the SDF function $\varphi_i : \mathcal{X} \mapsto \mathbb{R}$ is:

$$\varphi_i(\mathbf{x}) := \begin{cases} -d(\mathbf{x}, \partial \mathcal{O}_i), & \mathbf{x} \in \mathcal{O}_i, \\ d(\mathbf{x}, \partial \mathcal{O}_i), & \mathbf{x} \notin \mathcal{O}_i, \end{cases} \quad (2)$$

where d denotes the Euclidean distance between a point and a set. In this section, we describe an approach to construct approximations $\tilde{\varphi}_{k,i}$ to the obstacle SDFs φ_i using the point cloud observations $\{\mathcal{P}_{k,i}\}_k$ at time step t_k .

A. Data Pre-processing

Given the point cloud $\mathcal{P}_{k,i} \subset \mathbb{R}^n$ on the surface of the i th obstacle at time t_k , we can regard $\mathcal{P}_{k,i}$ as the points on the zero level-set of a distance function. To reduce the scale of training data, we define the origin of the i th obstacle $\bar{\mathbf{p}}_{k,i} := \arg \min_{\mathbf{p}} \sum_j \|\mathbf{p}_{k,i,j} - \mathbf{p}\|^2$. The points $\mathbf{p}_{k,i,j} - \bar{\mathbf{p}}_{k,i}$ are centered around $\bar{\mathbf{p}}_{k,i}$ and have a measured distance of 0 to the obstacle surface $\partial \mathcal{O}_i$. Let $c = \|\mathbf{x}_k - \mathbf{q}_{k,i,j}\|$ be the Euclidean distance between the robot state $\mathbf{x}_k := \mathbf{x}(t_k)$ and

$\mathbf{p}_{k,i,j}$. Let $\delta > 0$ be a small positive constant, define a point $\mathbf{q}_{k,i,j} := \frac{\delta}{c} \mathbf{x}_k + (1 - \frac{\delta}{c}) \mathbf{p}_{k,i,j}$ along the LiDAR ray from the robot state \mathbf{x}_k to $\mathbf{p}_{k,i,j}$ that is approximately a small distance δ from the obstacle surface $\partial\mathcal{O}_i$. We call the set $\{\mathbf{q}_{k,i,j} - \bar{\mathbf{p}}_{k,i}\}_j$ as truncated SDF point set.

For each obstacle i , the training set for the estimation task is constructed as a union of the points on the boundary and the truncated SDF points. The training set at time t_k is $\mathcal{D}_{k,i} := \{(\mathbf{p}_{k,i,j} - \bar{\mathbf{p}}_{k,i}, 0)\} \cup \{(\mathbf{q}_{k,i,j} - \bar{\mathbf{p}}_{k,i}, \delta)\}$.

B. Loss Function

Inspired by the recent impressive results on multilayer perceptron approximation of SDF [6], [8], we introduce a fully connected neural network $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_k) = \tilde{\varphi}_{k,i}$ with parameters $\boldsymbol{\theta}_k$ to approximate the SDF of each observed obstacle i at time step t_k . Our approach relies on the fact that an SDF φ has the property that its gradient norm satisfies the *Eikonal* equation $\|\nabla\varphi\| = 1$ in its domain. We use a loss function that encourages $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_k)$ to approximate the measured distances in a training set \mathcal{D} (distance loss ℓ_i^D) and to satisfy the Eikonal constraint for set of points \mathcal{D}' (Eikonal loss ℓ_i^E). For example, $\tilde{\varphi}_i$ equals 0 for points on the obstacle surface and equals δ for the truncated SDF points along the LiDAR rays. The loss function is defined as $\ell_i(\boldsymbol{\theta}_k; \mathcal{D}, \mathcal{D}') := \ell_i^D(\boldsymbol{\theta}_k; \mathcal{D}) + \lambda \ell_i^E(\boldsymbol{\theta}_k; \mathcal{D}')$, with a parameter $\lambda > 0$ and:

$$\begin{aligned} \ell_i^D(\boldsymbol{\theta}_k; \mathcal{D}) &:= \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{p}, d) \in \mathcal{D}} |\tilde{\varphi}_i(\mathbf{p}; \boldsymbol{\theta}_k) - d|, \\ \ell_i^E(\boldsymbol{\theta}_k; \mathcal{D}') &:= \frac{1}{|\mathcal{D}'|} \sum_{\mathbf{p} \in \mathcal{D}'} (\|\nabla \tilde{\varphi}_i(\mathbf{p}; \boldsymbol{\theta}_k)\| - 1). \end{aligned} \quad (3)$$

The Eikonal training set \mathcal{D}' can be generated arbitrarily in \mathcal{X} since $\tilde{\varphi}_i$ needs to satisfy the Eikonal constraint *everywhere*. In practice, we generate \mathcal{D}' by sampling points from a mixture of a uniform distribution and Gaussians centered at the points in the distance training set \mathcal{D} with standard deviation equal to the distance to the k -th nearest neighbor with $k = |\mathcal{D}|/2$. For the distance training set \mathcal{D} , ideally, we should use as much data as possible to obtain an accurate approximation of the SDF φ_i . For example, at time t_k , all observing data may be used as the distance training set $\mathcal{D} = \cup_{l=0}^k \mathcal{D}_{l,i}$. However, the online training time becomes longer and longer as the robot receives more and more LiDAR measurements, which makes it impractical for real-time mapping and navigation tasks. Therefore, we introduce an ‘‘Incremental Training with Replay Memory’’ approach in IV-C, which enables us to get accurate SDF estimates within reasonable training time for real-time tasks.

C. Incremental Learning

After training the network $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_0)$ when the obstacle i was first observed at time step $t = t_0$ with the training set $\mathcal{D}_{0,i}$, we need to consider how to update the MLP parameters $\boldsymbol{\theta}_k$ when we receive new observation sets $\mathcal{D}_{k,i}$ for $k = 1, 2, \dots$. We first consider an approach that updates $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_{k-1})$ based on the new observing set $\mathcal{D}_{k,i}$ and uses

the previous MLP parameters $\boldsymbol{\theta}_{k-1}$ as initialization, we call it ‘‘Incremental Training’’ (IT) approach. Alternatively, we consider updating $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_{k-1})$ based on all observing sets $\mathcal{D}_i = \cup_{l=0}^k \mathcal{D}_{l,i}$ up to time t_k and uses $\boldsymbol{\theta}_{k-1}$ as initialization. We call this update method as ‘‘Batch Training’’ (BT) approach.

The IT approach is efficient because at each time t_k , it uses training data sets $\mathcal{D}_{k,i}$ of approximately constant cardinality, leading to approximately constant update time during online training. However, discarding the old data $\cup_{l=0}^{k-1} \mathcal{D}_{l,i}$ and using stochastic gradient descent to re-train the network parameters $\boldsymbol{\theta}_{k-1}$ on the new data $\mathcal{D}_{k,i}$ causes degradation in the neural network’s ability to represent the old data. In other words, the SDF approximation $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_k)$ at time t_k is good at approximating the latest observed obstacle surface but the approximation quality degrades at previously observed surfaces. The BT approach does not have this limitation since it uses all training data $\cup_{l=0}^k \mathcal{D}_{l,i}$ for training at time t_k but, as a result, its training time increases over time. Hence, our motivation for introducing an ‘‘Incremental Training with Replay Memory’’ (ITRM) approach is to balance the trade-off between SDF estimation error and online training time, making it suitable for real-time robotic tasks.

Experience replay is an effective and commonly used technique for online reinforcement learning, which enables convergence of stochastic gradient descent to a critical point in policy and value function approximation [2], [12]. This idea has not been explored for online supervised learning of geometric surfaces. The first contribution of this paper is to use replay memory for online incremental learning of SDF. At each time step t_{k-1} , We construct an experience replay memory $\mathcal{Q}_{k-1,i}$ for each successive training step by utilizing the SDF approximations $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_{k-1})$.

Definition 3. The *replay memory* \mathcal{Q} associated with a signed distance function φ and truncation parameter $\tau \geq 0$ is a set of points that are at most a distance τ from the zero-level set of φ : $\mathcal{Q} := \{(\mathbf{q}, \varphi(\mathbf{q})) \in (\mathcal{X}, R) \mid |\varphi(\mathbf{q})| \leq \tau\}$.

To construct the replay memory set $\mathcal{Q}_{k-1,i}$, we need to generate samples from the level sets of the SDF approximation $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_{k-1})$. Matan *et al.* [13] introduced a simple and scalable approach to sample from neural network level sets. In robotics applications, where the environments are commonly 2-D or 3-D, samples from the level sets of $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_{k-1})$ can also be obtained using the Marching Cubes algorithm [14]. In our experiments in Sec. VI, we used the Marching Cubes algorithm to extract samples \mathbf{q}_0 and \mathbf{q}_δ from the zero and δ level-sets of $\tilde{\varphi}_i(\mathbf{x}; \boldsymbol{\theta}_{k-1})$, respectively, and construct the replay memory as $\mathcal{Q}_{k-1,i} := \{(\mathbf{q}_0, 0)\} \cup \{(\mathbf{q}_\delta, \delta)\}$.

Given the replay memory $\mathcal{Q}_{k-1,i}$, our ITRM approach constructs a training set at time t_k by combining the latest observation $\mathcal{D}_{k,i}$ with a randomly sampled subset $\bar{\mathcal{Q}}_{k-1,i}$ from $\mathcal{Q}_{k-1,i}$. To make the algorithm efficient without losing significant information about the previous data, we let $|\bar{\mathcal{Q}}_{k-1,i}| = |\mathcal{D}_{k,i}|$, which means that we pick the same number of points as the current observing set from the replay

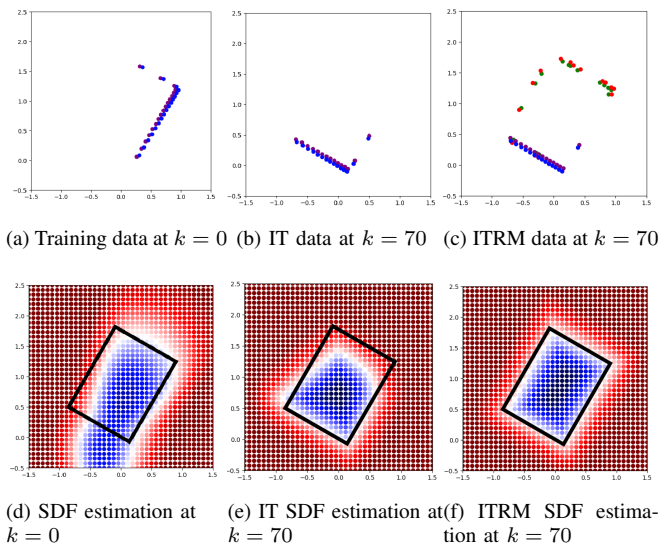


Fig. 2: Shape estimation with and without replay memory. The top row shows the training data used at time step $k = 0$ and $k = 70$ by the IT and ITRM approaches. The purple points are the observed LiDAR end points, while the blue points are the truncated SDF points along the LiDAR rays. In (c), the green and red points are boundary and truncated SDF points obtained from the replay memory. The bottom row shows the SDF estimation of the obstacle surface at time step $k = 0$ and $k = 70$ for the IT and ITRM approaches. The black rectangle shows the ground-truth obstacle boundary, while the colored regions are the level-sets of the SDF estimate. The white region denotes the estimated obstacle boundary. The blue region denotes negative signed distance (inside the obstacle). The red region denotes positive signed distance (outside the obstacle). Note that IT and ITRM use the same data and lead to the same estimate at $k = 0$ because the replay memory set is empty. In (e), the SDF estimate of the top obstacle region at $k = 70$, without memory replay, degrades compared to (d). In (f), training with replay memory, helps the neural network remember the overall obstacle shape.

memory.

The three neural network training techniques, IT, BT, and ITRM, discussed in this section, are formally defined as follows. In all approaches, for obstacle i , the MLP parameters θ_{k-1} obtained at the previous time step are used as an initial guess for the new MLP parameters θ_k at time t_k . The SDF approximations $\tilde{\varphi}_i(\mathbf{x}; \theta_k)$ are trained via the loss function in (3) but each method uses a different training set \mathcal{D} :

- IT: $\mathcal{D} = \mathcal{D}_{k,i}$,
- BT: $\mathcal{D} = \bigcup_{l=0}^k \mathcal{D}_{l,i}$,
- ITRM: $\mathcal{D} = \mathcal{D}_{k,i} \cup \bar{\mathcal{Q}}_{k,i}$.

We use Softplus (SmoothReLU function) as the non-linear activations of our network, which is defined as $w(x) = \ln(1 + e^x)$. As the network is a combination of linear functions and Softplus, the output $\tilde{\varphi}_i(\mathbf{x}; \theta_k)$ is a continuously differentiable function. The estimated gradient $\nabla \tilde{\varphi}_i(\mathbf{x}; \theta_k)$ is one input in our loss function and can be calculated via backpropagation through the neural network layers [8].

Remark 1 (Errors in the network approximation). We do not formally compute the error bounds in the SDF approximation

and leave the theoretical computation of SDF approximation error for future work. Several recent works study the approximation power and error bounds of neural networks [15], [16]. Our ensuing analysis does take into account errors in this approximation to synthesize safe controllers. In our evaluation, cf. Sec. VI, we obtain the SDF error bounds by comparing the SDF estimations to the ground-truth object SDFs, and use them to help perform online navigation tasks.

V. SAFE NAVIGATION WITH ESTIMATED OBSTACLES

In this section, we rely on the estimated SDFs constructed in Sec. IV to formalize the synthesis of a controller that guarantees safety with respect to the exact obstacles, despite the errors in the approximation.

A. Safe Control with Estimated Barrier Functions

A useful tool to ensure that the robot state remains in the safe set \mathcal{S} throughout its evolution is the notion of control barrier function (CBF).

Definition 4 ([17]). A continuously differentiable function $h : \mathbb{R}^n \mapsto \mathbb{R}$, with $h(\mathbf{x}) > 0$ if $\mathbf{x} \in \text{Int}(\mathcal{S})$ and $h(\mathbf{x}) = 0$ if $\mathbf{x} \in \partial\mathcal{S}$, is a *zeroing control barrier function (ZCBF)* on $\mathcal{X} \subset \mathbb{R}^n$ if there exists an extended class \mathcal{K} function α_h such that, for each $\mathbf{x} \in \mathcal{X}$, there exists $\mathbf{u} \in \mathcal{U}$ with

$$\mathcal{L}_f h(\mathbf{x}) + \mathcal{L}_g h(\mathbf{x})\mathbf{u} + \alpha_h(h(\mathbf{x})) \geq 0, \quad (4)$$

where $\mathcal{L}_f h(\mathbf{x})$ is the Lie derivative of $h(\mathbf{x})$ along $f(\mathbf{x})$.

Any Lipschitz-continuous controller $\mathbf{u} : \mathcal{X} \mapsto \mathcal{U}$ such that $\mathbf{u}(\mathbf{x})$ satisfies (4) renders the set \mathcal{S} forward invariant for the control-affine system (1), cf [1], [17].

If the exact SDFs φ_i describing the obstacles \mathcal{O}_i were known, they could be used to define ZCBFs $h_i(\mathbf{x}) := \varphi_i(\mathbf{x})$ that ensure the forward invariance of \mathcal{S} . However, when the environment is observable only through online range measurements as described in Sec. IV, we only have the estimated SDFs $\tilde{\varphi}_i(\mathbf{x}; \theta_k)$ at our disposal to define $\tilde{h}_i(\mathbf{x}) := \tilde{\varphi}_i(\mathbf{x}; \theta_k)$ (note that \tilde{h}_i is continuously differentiable). Our next result describes how to use this information to ensure the safety of \mathcal{S} . The statement is for a generic \tilde{h} (e.g., a single obstacle) and we later particularize our discussion to the case of multiple obstacles.

Proposition 1. Let $e_1(\mathbf{x}) := h(\mathbf{x}) - \tilde{h}(\mathbf{x}) \in \mathbb{R}$ be the error at \mathbf{x} in the approximation of h , and likewise, let $e_2(\mathbf{x}) := \nabla h(\mathbf{x}) - \nabla \tilde{h}(\mathbf{x}) \in \mathbb{R}^n$ be the error at \mathbf{x} in the approximation of its gradient. Assume there are available known functions $e_h(\mathbf{x}) : \mathbb{R} \mapsto \mathbb{R}_{\geq 0}$ and $e_{\nabla h}(\mathbf{x}) : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}$ such that

$$|e_1(\mathbf{x})| \leq e_h(\mathbf{x}), \quad \|e_2(\mathbf{x})\| \leq e_{\nabla h}(\mathbf{x}) \quad (5)$$

with $e_h(\mathbf{x}) \rightarrow 0$ and $e_{\nabla h}(\mathbf{x}) \rightarrow 0$ as $\mathbf{x} \rightarrow \partial\mathcal{S}$. Let

$$\mathcal{K}_{\tilde{h}}(\mathbf{x}) := \{\mathbf{u} \in \mathcal{U} \mid \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} - \|f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}\| e_{\nabla h}(\mathbf{x}) + \alpha_h(\tilde{h}(\mathbf{x}) - e_h(\mathbf{x})) \geq 0\}. \quad (6)$$

Then, any locally Lipschitz continuous controller $\mathbf{u} : \mathcal{X} \mapsto \mathcal{U}$ such that $\mathbf{u}(\mathbf{x}) \in \mathcal{K}_{\tilde{h}}(\mathbf{x})$ guarantees that the safe set \mathcal{S} is forward invariant.

Proof. We start by substituting $h(\mathbf{x}) = \tilde{h}(\mathbf{x}) + e_1(\mathbf{x})$ in (4):

$$\begin{aligned} \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} + e_2(\mathbf{x})^\top f(\mathbf{x}) + e_2(\mathbf{x})^\top g(\mathbf{x})\mathbf{u} \\ \geq -\alpha_h(\tilde{h}(\mathbf{x}) + e_1(\mathbf{x})). \end{aligned}$$

For any fixed \mathbf{x} and any errors $e_1(\mathbf{x})$ and $e_2(\mathbf{x})$ satisfying (5), we need the minimum value of the left-hand side greater than the maximum value of the right-hand side to ensure that (4) still holds, namely:

$$\min_{\|e_2(\mathbf{x})\| \leq e_{\nabla h}(\mathbf{x})} \left\{ \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} + e_2(\mathbf{x})^\top f(\mathbf{x}) + e_2(\mathbf{x})^\top g(\mathbf{x})\mathbf{u} \right\} \geq \max_{|e_1(\mathbf{x})| \leq e_h(\mathbf{x})} \left\{ -\alpha_h(\tilde{h}(\mathbf{x}) + e_1(\mathbf{x})) \right\}. \quad (7)$$

Note that since $e_h(\mathbf{x}) \geq 0$ and α_h is an extended class \mathcal{K} function, the maximum value in (7) is obtained by:

$$\max_{|e_1(\mathbf{x})| \leq e_h(\mathbf{x})} -\alpha_h(\tilde{h}(\mathbf{x}) + e_1(\mathbf{x})) = -\alpha_h(\tilde{h}(\mathbf{x}) - e_h(\mathbf{x})).$$

The minimum value is attained when $e_2(\mathbf{x})$ is in the opposite direction to the gradient of $f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}$, namely

$$\min_{\|e_2(\mathbf{x})\| \leq e_{\nabla h}(\mathbf{x})} \left\{ e_2(\mathbf{x})^\top f(\mathbf{x}) + e_2(\mathbf{x})^\top g(\mathbf{x})\mathbf{u} \right\} = -\|f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}\| e_{\nabla h}(\mathbf{x}).$$

Therefore, the inequality condition (7) can be rewritten as

$$\begin{aligned} \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} - \|f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}\| e_{\nabla h}(\mathbf{x}) \\ \geq -\alpha_h(\tilde{h}(\mathbf{x}) - e_h(\mathbf{x})), \end{aligned}$$

which is equivalent to the condition in (6). \square

Proposition 1 allows us to synthesize safe controllers even though the obstacles are not exactly known, provided that error bounds on the approximation of the barrier function and its gradient are available and get better as the robot state gets closer to the boundary of the safe set \mathcal{S} .

B. Control Synthesis via Second-Order Cone Programming

We encode the control objective using the notion of a Lyapunov function. Formally, we assume the existence of a control Lyapunov function, as defined next.

Definition 5 ([1]). A *control Lyapunov function (CLF)* for the system dynamics in (1) is a continuously differentiable function $V : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}$ for which there exist a class \mathcal{K} function α_V such that, for all $\mathbf{x} \in \mathcal{X}$:

$$\inf_{\mathbf{u} \in \mathcal{U}} [\mathcal{L}_f V(\mathbf{x}) + \mathcal{L}_g V(\mathbf{x})\mathbf{u} + \alpha_V(V(\mathbf{x}))] \leq 0.$$

The function V may be used to encode a variety of control objectives, including for instance path following. We present a specific Lyapunov function for this purpose in Sec.VI.

When the barrier function is precisely known, one can combine CLF and CBF constraints to synthesize a safe controller via the following Quadratic Program:

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \delta \in \mathbb{R}} \|L(\mathbf{x})^\top (\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x}))\|^2 + \lambda \delta^2 \\ \text{s.t. } \mathcal{L}_f V(\mathbf{x}) + \mathcal{L}_g V(\mathbf{x})\mathbf{u} + \alpha_V(V(\mathbf{x})) \leq \delta \\ \mathcal{L}_f h(\mathbf{x}) + \mathcal{L}_g h(\mathbf{x})\mathbf{u} + \alpha_h(h(\mathbf{x})) \geq 0, \end{aligned} \quad (8)$$

where $\tilde{\mathbf{u}}(\mathbf{x})$ is a baseline controller, $L(\mathbf{x})$ is a matrix penalizing control effort, and $\delta \geq 0$ (with the corresponding penalty λ) is a slack variable, introduced to relax the CLF constraints in order to ensure the feasibility of the QP. The baseline controller $\tilde{\mathbf{u}}(\mathbf{x})$ is used to specify additional control requirements such as desirable velocity or orientation (see Sec. VI) but may be set to $\tilde{\mathbf{u}}(\mathbf{x}) \equiv 0$ if minimum control effort is the main objective. The QP formulation in (8) modifies $\tilde{\mathbf{u}}(\mathbf{x})$ online to ensure safety and liveness via the CBF and CLF constraints.

Without exact knowledge of the barrier function h , in order to synthesize a safe controller via (8), we need to replace (4) by (6):

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \delta \in \mathbb{R}} \|L(\mathbf{x})^\top (\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x}))\|^2 + \lambda \delta^2 \\ \text{s.t. } \mathcal{L}_f V(\mathbf{x}) + \mathcal{L}_g V(\mathbf{x})\mathbf{u} + \alpha_V(V(\mathbf{x})) \leq \delta \\ \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} + \alpha_h(\tilde{h}(\mathbf{x}) - e_h(\mathbf{x})) \\ \geq \|f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}\| e_{\nabla h}(\mathbf{x}). \end{aligned} \quad (9)$$

This makes the optimization problem in (9) no longer a quadratic program. However, the following result shows that (9) is a second-order cone program (SOCP), and can therefore be solved efficiently.

Proposition 2. *The optimization problem in (9) is equivalent to the following second-order cone program:*

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \delta \in \mathbb{R}, l \in \mathbb{R}} l \\ \text{s.t. } \mathcal{L}_f V(\mathbf{x}) + \mathcal{L}_g V(\mathbf{x})\mathbf{u} + \alpha_V(V(\mathbf{x})) \leq \delta \\ \|f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}\| e_{\nabla h}(\mathbf{x}) \leq \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} \\ + \alpha_h(\tilde{h}(\mathbf{x}) - e_h(\mathbf{x})) \\ \left\| \begin{bmatrix} 2L(\mathbf{x})^\top (\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x})) \\ 2\sqrt{\lambda}\delta \\ l - 1 \end{bmatrix} \right\| \leq l + 1. \end{aligned} \quad (10)$$

Proof. We first introduce a new variable l so that the problem in (9) is equivalent to

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \delta \in \mathbb{R}, l \in \mathbb{R}} l \\ \text{s.t. } \mathcal{L}_f V(\mathbf{x}) + \mathcal{L}_g V(\mathbf{x})\mathbf{u} + \alpha_V(V(\mathbf{x})) \leq \delta \\ \|f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}\| e_{\nabla h}(\mathbf{x}) \leq \mathcal{L}_f \tilde{h}(\mathbf{x}) + \mathcal{L}_g \tilde{h}(\mathbf{x})\mathbf{u} \\ + \alpha_h(\tilde{h}(\mathbf{x}) - e_h(\mathbf{x})) \\ \|L(\mathbf{x})^\top (\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x}))\|^2 + \lambda \delta^2 \leq l. \end{aligned} \quad (11)$$

The last constraint in (11) corresponds to a rotated second-order cone, $\mathcal{Q}_{rot}^n := \{(\mathbf{x}_r, y_r, z_r) \in \mathbb{R}^{n+2} \mid \|\mathbf{x}_r\|^2 \leq y_r z_r, y_r \geq 0, z_r \geq 0\}$, which can be converted into a standard SOC constraint [18]:

$$\left\| \begin{bmatrix} 2\mathbf{x}_r \\ y_r - z_r \end{bmatrix} \right\| \leq y_r + z_r.$$

Let $y_r = l$, $z_r = 1$ and consider the constraint $\|L(\mathbf{x})^\top (\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x}))\|^2 + \lambda \delta^2 \leq l$. Multiplying both sides by 4 and adding $l^2 + 1$, the constraint is equivalent to

$$4\|L(\mathbf{x})^\top (\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x}))\|^2 + 4\lambda \delta^2 + (l - 1)^2 \leq (l + 1)^2.$$

Taking a square root on both sides, we end up with $\sqrt{\|2L(\mathbf{x})^\top(\mathbf{u} - \tilde{\mathbf{u}}(\mathbf{x}))\|^2 + (2\sqrt{\lambda}\delta)^2 + (l-1)^2} \leq l+1$, which is equivalent to the third constraint in (10). \square

For multiple obstacles in the environment, one can add multiple CBF constraints to (10). We leave for future work the characterization of the Lipschitz continuity properties of the controller resulting from (10).

VI. EXPERIMENTS AND RESULTS

Here we use the Pybullet simulator [19] to evaluate the performance of our proposed approach for online obstacle shape estimation and navigation with robust safety constraints. In our experiments, we use a two-wheeled ground TurtleBot, equipped with a LiDAR scanner, which has a 270 degree field of view and 3 meters maximum range. The simulation environments contain obstacles with various shapes that are unknown to the robot. The goal is to show that the robot can avoid obstacles with the estimated barrier functions while following the desired path whenever possible.

A. Modeling

We model the robot motion using unicycle kinematics and take a small distance $a \neq 0$ off the wheel axis as in [20] to obtain a relative-degree-one model:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} v \cos(\theta) - a\omega \sin(\theta) \\ v \sin(\theta) + a\omega \cos(\theta) \\ \omega \end{bmatrix}, \quad (12)$$

where v , ω represent the robot longitudinal and angular velocity, respectively. The state and input are $\mathbf{x} := [x, y, \theta]^\top \in \mathbb{R}^2 \times [-\pi, \pi)$, $\mathbf{u} := [v, \omega]^\top \in \mathbb{R}^2$. We use a baseline controller $\tilde{\mathbf{u}}(\mathbf{x}) \equiv [v_{max}, 0]^\top$ to encourage the robot to drive at its maximum velocity v_{max} in a straight line whenever possible.

We specify the desired robot path by a 2-D curve $r(\gamma)$, $\gamma \in [0, 1]$. To capture the path-following task via a CLF constraint, we define an output function $\boldsymbol{\eta}(\mathbf{x}) := \mathbf{x} - r(\gamma(\mathbf{x}))$. The closest point on the reference path $r(\gamma)$ to the robot state \mathbf{x} is obtained by $\gamma(\mathbf{x}) := \arctan(y/x)/(2\pi) \in [0, 1]$. According to the results in [21], the following function is a valid CLF for the path following task:

$$V(\mathbf{x}) = [\boldsymbol{\eta}^\top(\mathbf{x}), \dot{\boldsymbol{\eta}}^\top(\mathbf{x})]P[\boldsymbol{\eta}^\top(\mathbf{x}), \dot{\boldsymbol{\eta}}^\top(\mathbf{x})]^\top, \quad (13)$$

where P is a positive-definite matrix calculated by solving the Lyapunov equation of the input-output linearization.

B. Network Architecture and Training Implementation

We used the MLP network architecture proposed by Gropp *et al.* [8] and Park *et al.* [6] to represent the SDF approximator $\tilde{\varphi}(\mathbf{x}; \boldsymbol{\theta}_k)$. The neural network has 8 fully connected layers with 512 neurons each and a single skip connection from the input to the middle layer. All internal layers use Softplus as nonlinear activations. We set the parameter $\lambda = 0.1$ in (3). For online incremental training with replay memory, at each time step t_k , we performed 16 iterations with the ADAM optimizer [22] with constant learning rate of 0.001.

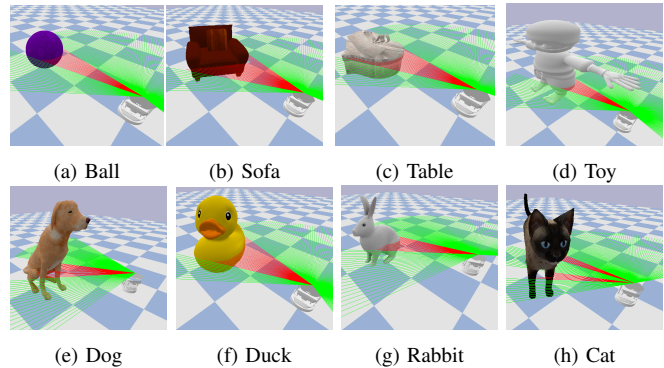


Fig. 3: Object instances used to evaluate the estimation performance of online signed distance function approximation.

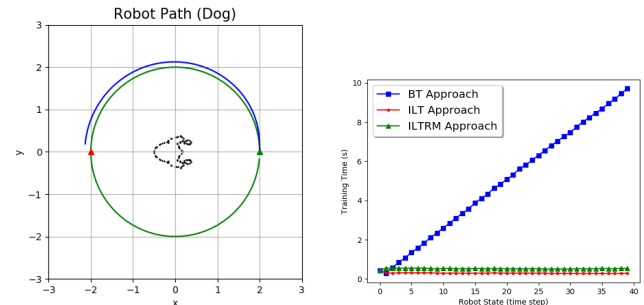


Fig. 4: Robot motion (left) and training time results (right) for the SDF estimation experiment in Sec. VI-C. The robot starts at $(2, 0)$ and follows a circular reference path (green) to a goal position at $(-2, 0)$ (red triangle). The actual trajectory followed by the controller is shown in blue. The black dots are points on the surface of a ground-truth Dog object and vary for different object instance. The average training time for the BT, ILT, and ILTRM methods for SDF approximation is compared across the 8 object instances shown in Fig. 3.

C. SDF Estimation Results

To evaluate the accuracy and efficiency of the proposed ITRM approach for SDF estimation, we experimented with various obstacle shapes and compare the SDF prediction accuracy and the training time of our proposed ITRM approach and the IT and BT approaches described in Sec. IV-C. Fig. 3 and Fig. 4 show our experiment setup. The robot is equipped with a LiDAR with 150 rays per scan. Zero-mean Gaussian noise with standard deviation $\sigma = 0.01$ was added to the distance measurements along each ray. To evaluate the accuracy of the SDF approximation $\tilde{\varphi}(\mathbf{x}_i; \boldsymbol{\theta}_k)$, we sample $m = 200$ points $\{\mathbf{x}_i\}_{i=1}^m$ uniform on the surface of the ground-truth object and measure the SDF error (k varies for different objects):

$$\mathcal{E} = \frac{1}{m} \sum_{i=1}^m |\tilde{\varphi}(\mathbf{x}_i; \boldsymbol{\theta}_k)|. \quad (14)$$

The SDF error is shown in Table I for the eight object instance in Fig. 3. The SDF error of our ITRM approach is comparable with the error of BT approach and is much smaller than the error of the IT approach across the 8 object instances. The training update time is the time needed for

TABLE I: SDF estimation error (14) of the IT, ITRM, and BT training methods for LiDAR measurements with zero-mean Gaussian range noise with standard deviation $\sigma = 0.01$.

Case	\mathcal{E} in IT	\mathcal{E} in ITRM	\mathcal{E} in BT
1 (Ball)	0.028722	0.014783	0.013994
2 (Sofa)	0.065302	0.021212	0.021509
3 (Table)	0.056442	0.023013	0.022738
4 (Toy)	0.106427	0.024153	0.019792
5 (Dog)	0.032874	0.021165	0.012306
6 (Duck)	0.075069	0.010226	0.011129
7 (Rabbit)	0.060345	0.013431	0.011008
8 (Cat)	0.029232	0.014867	0.012262
Average	0.056802	0.017856	0.014691

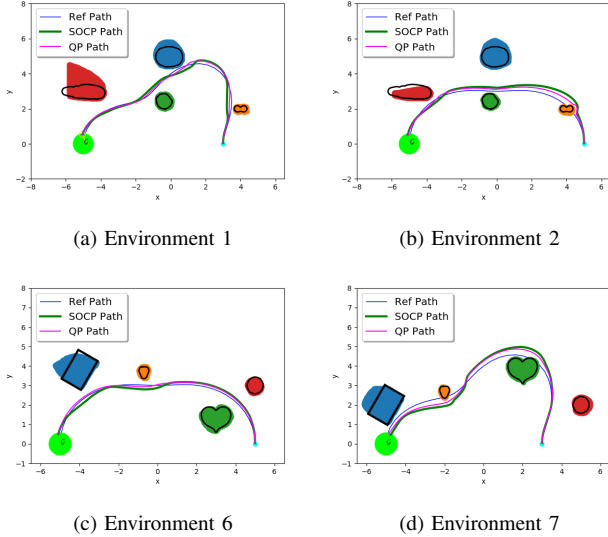


Fig. 5: Simulation results for the CLF-CBF-SOCP and CLF-CBF-QP frameworks. The ground-truth obstacle surfaces are drawn in black. The estimated obstacles, obtained after the whole trajectory is traversed, under CLF-CBF-SOCP framework are shown in different colors (red, green, blue, orange). The reference path is shown in blue. The robot trajectory generated by CLF-CBF-SOCP is shown in green, while the trajectory generated by CLF-CBF-QP is shown in pink. The starting point is the cyan point and the goal region is a light green circle.

updating the MLP parameters from θ_{k-1} to θ_k . The average training update time across the 8 object instances for the three methods is shown in Fig. 4. We see that as the robot moves around the environment, the average training update time of the ITRM approach remains at about $0.7s$ while the BT approach requires more and more time for training.

D. Safe Navigation Results

The second set of experiments demonstrates safe trajectory tracking using online SDF obstacle estimates to define constraints in the CLF-CBF-SOCP control synthesis optimization (10). The robot is asked to move along a reference path while avoiding a priori unknown obstacles in 8 different environments, similar to the one shown in Fig. 1. To account for the fact that the robot body is not a point mass, we subtract the robot radius $\gamma = 0.177$ from the SDF estimate when defining the CBF: $h_i(\mathbf{x}) = \tilde{\varphi}_i(\mathbf{x}; \theta) - \gamma$. The error bounds of CBF $e_h(\mathbf{x})$ and its gradient $e_{\nabla h}(\mathbf{x})$ are approx-

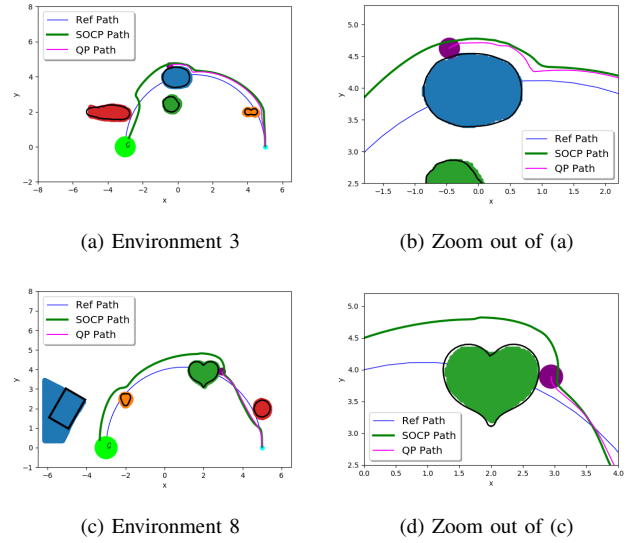


Fig. 6: Simulation results in environments where the CLF-CBF-SOCP controller succeeded but CLF-CBF-QP one failed. The robot is shown as purple circle when crashing into an obstacle using the CLF-CBF-QP controller. Fig. 6b and Fig. 6d show enlargements of the crashing regions in Fig. 6a and Fig. 6c, respectively.

imated based on Table I. We compare the performance of the proposed CLF-CBF-SOCP approach to a CLF-CBF-QP approach that assumes the estimated obstacle CBFs $\hat{h}_i(\mathbf{x})$ are accurate and ignores the estimation error. In the experiment, to avoid small linear velocity and ensure the feasibility of our problem, we set the relaxation penalty parameter on the diagonal of $L(\mathbf{x})$ in Sec. V-B as $l_1 = 10, l_2 = 1, l_3 = 10\sqrt{10}$ at the beginning and make l_1 adaptive. l_1, l_2, l_3 are the penalty parameters for linear velocity, angular velocity and path following respectively. If there is no solution found at some time step, l_1 will be divided by $\sqrt{2}$ until one feasible solution found.

We use the Fréchet distance between the reference path r and the paths produced by the CLF-CBF-SOCP and CLF-CBF-QP frameworks to evaluate the path similarities. Let A, B be two paths in the environment as defined in Def.1, the Fréchet distance $F(A, B)$ is a measure of similarity between A and B that takes into account the location and ordering of the points along the paths. Two reparameterizations α, β of $[0, 1]$ are continuous, non-decreasing, surjections $\alpha, \beta : [0, 1] \mapsto [0, 1]$. The Fréchet distance is computed by:

$$F(A, B) = \inf_{\alpha, \beta} \max_{t \in [0, 1]} \{d(A(\alpha(t)), B(\beta(t)))\}, \quad (15)$$

where d is the Euclidean distance in our case.

In Fig. 5, we see that the robot can follow the prescribed reference path if no obstacles are nearby. Comparing the green and pink trajectories in the 4 plots, we see that the green trajectory generated by CLF-CBF-SOCP is always more conservative since it takes the error in the CBF estimates into account. When there is an obstacle near or on the reference path, the robot controlled by the CLF-CBF-SOCP controller stays further away from the obstacles than

TABLE II: Fréchet distance between the reference path and the robot trajectories generated by the CLF-CBF-SOCP and the CLF-CBF-QP controllers (smaller values indicate larger trajectory similarity, N\A indicates that the robot failed to reach the goal region)

Environment	Online CLF-CBF-SOCP vs Groundtruth	Online CLF-CBF-QP vs Groundtruth
1	0.342766	0.344171
2	0.533891	0.397145
3	0.835619	N\A
4	0.396763	0.349201
5	1.053466	N\A
6	0.361139	0.344407
7	0.464251	0.349328
8	0.867368	N\A

the robot controlled by the CLF-CBF-QP controller. This agrees with the Fréchet distance results presented in Table II.

In Fig. 6, we see that the CLF-CBF-QP controller sometimes fails to avoid obstacles because it does not take the error in the CBF estimation into account. In contrast, the CLF-CBF-SOCP controller is guaranteed by Prop. 2 to remain safe if the CBF estimation is captured correctly in the SOC constraints.

To sum up, from Table II, we notice that the trajectory mismatch with respect to the reference path is larger under our proposed CLF-CBF-SOCP framework than under the CLF-CBF-QP framework if they both succeed. However, our proposed approach guarantees safe navigation while the CLF-CBF-QP controller sometimes leads to collision due to errors in SDF estimation.

VII. CONCLUSION

We have proposed the Incremental Training with Replay Memory method to approximate control barrier functions. We introduced the first online incremental learning approach in estimating Signed Distance Functions that utilizes replay memory. By using replay memory in training, our proposed approach balances the estimation error and the training time, which ensures it suitable for real-time robotic tasks. The CBF approximations are then used to formulate a novel CLF-CBF-SOCP framework which ensures safe navigation in an unknown environment by taking measurement and estimation uncertainties into account to synthesize a robust navigation algorithm. Compared with the traditional CLF-CBF-QP framework, our proposed approach ensures safety in arbitrary environments.

REFERENCES

- [1] A. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *European Control Conference (ECC)*, pp. 3420–3431, 2019.
- [2] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," in *International Conference on Learning Representations (ICLR)*, 2016.
- [3] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, 2013.
- [4] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1366–1373, 2017.

- [5] L. Han, F. Gao, B. Zhou, and S. Shen, "Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [6] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [7] L. M. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4455–4465, 2019.
- [8] A. Gropp, L. Yariv, N. Haim, M. Atzmon, and Y. Lipman, "Implicit geometric regularization for learning shapes," in *International Conference on Machine Learning*, 2020.
- [9] P. Glotfelter, J. Cortés, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," *IEEE Control Systems Letters*, vol. 1, no. 2, pp. 310–315, 2017.
- [10] X. Xu, T. Waters, D. Pickem, P. Glotfelter, M. Egerstedt, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Realizing simultaneous lane keeping and adaptive speed regulation on accessible mobile robot testbeds," in *IEEE Conference on Control Technology and Applications (CTA)*, pp. 1769–1775, 2017.
- [11] M. Srinivasan, A. Dabholkar, S. Coogan, and P. Vela, "Synthesis of control barrier functions using a supervised machine learning approach," *arXiv preprint:2003.04950*, 2020.
- [12] D. Rolnick, A. Ahuja, J. Schwarz, T. Lillicrap, and G. Wayne, "Experience replay for continual learning," in *AAAI Conference on Artificial Intelligence*, 2018.
- [13] M. Atzmon, N. Haim, L. Yariv, O. Israelov, H. Maron, and Y. Lipman, "Controlling neural level sets," in *Advances in Neural Information Processing Systems*, pp. 2032–2041, 2019.
- [14] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," *COMPUTER GRAPHICS*, vol. 21, no. 4, pp. 163–169, 1987.
- [15] B. Bailey, Z. Ji, M. Telgarsky, and R. Xian, "Approximation power of random neural networks," *ArXiv*, vol. abs/1906.07709, 2019.
- [16] D. Yarotsky, "Error bounds for approximations with deep relu networks," *Neural Networks*, vol. 94, pp. 103–114, 2017.
- [17] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [18] F. Alizadeh and D. Goldfarb, "Second-order cone programming," *Mathematical programming*, vol. 95, no. 1, pp. 3–51, 2003.
- [19] E. Coumans and Y. Bai, "PyBullet, a Python module for physics simulation for games, robotics and machine learning." <http://pybullet.org>, 2016.
- [20] J. Cortés and M. Egerstedt, "Coordinated control of multi-robot systems: A survey," *SICE Journal of Control, Measurement, and System Integration*, vol. 10, no. 6, pp. 495–503, 2017.
- [21] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle, "Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics," *IEEE Transactions on Automatic Control*, vol. 59, no. 4, pp. 876–891, 2014.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint: 1412.6980*, 2015.