Learning to Nudge: A Scalable Barrier Function Framework for Safe Robot Interaction in Dense Clutter

Haixin Jin*,[1], Nikhil Uday Shinde*,[1], Soofiyan Atar[1], Hongzhan Yu[1], Dylan Hirsch[1],
Sicun Gao[1], Michael C. Yip[1], Sylvia Herbert[1]
[1]University of California San Diego
{haj013, nshinde, satar, hoy021, dhirsch, sig049, yip, sherbert}@ucsd.edu

***Abstract*— Robots operating in everyday environments must navigate and manipulate within densely cluttered spaces, where physical contact with surrounding objects is unavoidable. Traditional safety frameworks treat contact as unsafe, restricting robots to collision avoidance and limiting their ability to function in dense, everyday settings. As the number of objects grows, model-based approaches for safe manipulation become computationally intractable; meanwhile, learned methods typically tie safety to the task at hand, making them hard to transfer to new tasks without retraining. In this work we introduce Dense Contact Barrier Functions(DCBF). Our approach bypasses the computational complexity of explicitly modeling multi-object dynamics by instead learning a composable, object-centric function that implicitly captures the safety constraints arising from physical interactions. Trained offline on interactions with a few objects, the learned DCBFcomposes across arbitrary object sets at runtime, producing a single global safety filter that scales linearly and transfers across tasks without retraining. We validate our approach through simulated and hardware experiments in dense clutter, demonstrating its ability to enable safe, contact-rich interaction in suitable settings.**

I. Introduction

The ability for robots to safely interact with and maneuver through densely cluttered environments is a critical step towards their widespread deployment in everyday, unstructured spaces. This could enable household robots to safely navigate crowded pantries, or allow warehouse robots to restock shelves in tight commercial settings. These settings contain dense, dynamic arrangements of objects, forcing the robot to make physical contact to complete its tasks. For example, a robot retrieving an item from a pantry may need to nudge aside items, while a logistics robot may need to make contact with fragile goods without knocking them over.

Research on robot safety has historically focused on collision avoidance. The majority of previous work from classical methods like Hamilton-Jacobi Reachability (HJR) [1], Control Barrier Functions (CBFs) [2] and Artificial Potential Fields (APF) [3], which filter a nominal control to output a safe control action, are all predicated on maintaining some minimum distance from all obstacles. Even works focused on complex manipulation often handle safety by incorporating the manipulated object into the set of items with which collision must be avoided [4], [5]. However, this criterion of collision avoidance is fundamentally limiting in cluttered spaces. A home robot focused on strictly avoiding collisions would be paralyzed with the task of reaching any object beyond the first row of items in a pantry. Humans, in contrast, routinely make contact, nudging and shifting objects while implicitly understanding the interactions required to prevent spills or damage. This work aims to bridge this gap, giving robots the capacity to reason about and execute safe contact.

While some research has explored physical interaction, these methods often fail to address the scalability and generalizability required for these dense multi-object scenarios. Approaches that explicitly model the complex, coupled dynamics between the robot and object [6] become computationally intractable as the number of objects grows, due to the combinatorial explosion of potential interactions. Works such as [6], [7] demonstrate safe contact with a single deformable object, but fail to provide a clear path to scaling to multi-object systems. Learned methods, such as those in reinforcement learning [8], [9], can handle complex systems but entangle safety and performance, requiring complete retraining for each new task. This highlights the need for a method that balances performance with safety during these interactions, in a manner that scales to a large, variable number of objects and generalizes to diverse arrangements.

In this paper we introduce Dense Contact Barrier Functions (DCBF), object-centric control barrier functions that implicitly embed learned interaction effects into the CBF landscape to achieve safe, scalable interaction in dense object environments. The core of our framework is a composable DCBF that is trained offline on interactions with just a few objects. At runtime, this function is evaluated for each object in the scene, and the results are aggregated into a single safety filter. This compositional approach allows our method to scale to environments with an arbitrary number of objects and generalizes across diverse configurations without retraining. Our method avoids the intractability of explicit multi-object interaction modeling and reasoning by learning to implicitly capture the safety constraints arising from complex interactions in an object-centric, scalable manner. Our primary contributions are:

- An implicit interaction CBF formulation that captures the multi-object interaction effects on safety without explicit modeling.
- A composable, object-centric framework to enable scalable safety reasoning for a variable number of objects.
- A high-interaction refinement procedure to reduce the

*These authors contributed equally to this work.

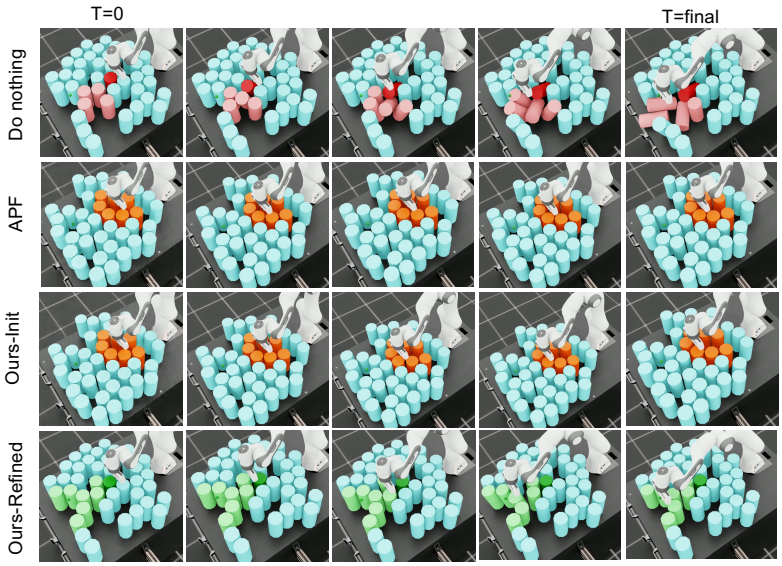

Fig. 1: Simulated Robot trajectory from start to goal across different methods. The first row shows the *Do Nothing* baseline, where the robot directly pushes through the clutter, causing multiple objects to be knocked over. Interacted objects are highlighted in red, with the object starting closest to the end-effector shown in dark red and others in light red. The second row shows the *Artificial Potential Field* baseline. The end-effector becomes stuck in clutter and fails to reach the goal. The objects responsible for the blockage are highlighted in orange. The third row shows the *Init Model* baseline. Similarly, the end-effector also gets stuck and cannot complete the task, with the blocking objects highlighted in orange. The fourth row shows our proposed method, which safely interacts with objects to reach the goal without violating safety. Interacted objects are highlighted in green, with the starting closest object shown in dark green and others in light green.

conservativeness of the learned CBF
- Validation of our approach in simulated experiments for safe interaction in densely cluttered scenes.

II. Related Works

The development of safe autonomous systems is a critical area of research, particularly for robots operating in complex dynamic environments such as households or commercial facilities. While a large body of work exists on ensuring robot safety, much of it has centered on collision avoidance. This paradigm, though effective in open spaces, is overly restrictive in dense settings where robots must make physical contact and interact with objects to complete their tasks. Our work addresses this gap by proposing a scalable method for safe robot–object interaction through composable barrier functions and implicit interaction modeling.

Safe Autonomy for Collision Avoidance: Robot safety has traditionally been synonymous with avoiding collisions. Classical methods like Hamilton-Jacobi Reachability (HJR) compute safety value functions that define control invariant safe sets [1]. These functions have been used online for optimal control and as minimally invasive safety filters [10]. However, HJR suffers from the curse of dimensionality and requires analytical models, limiting its use to low dimensional, control affine systems and consequently simpler collision avoidance based safety. Recent works like [11], [12] extend HJR using learning-based approximations, alleviating dimensionality restrictions but they still depend on known analytic models, limiting their scalability to environments with unknown dynamic object interactions.

This focus on collision avoidance persists even in manipulation tasks where object interaction is the goal [13]. Classical methods such as potential fields [3], [14] and MPC-based planners [15] have been focused on maintaining separation from and minimal contact with objects. More recent work in cluttered [5] and deformable object manipulation [4] simply extends this notion of safety by treating the manipulated object as part of the system, thus ensuring the object also avoids collisions with the environment. These approaches remain overly conservative for crowded environments where avoiding contact makes tasks infeasible.

Challenges in Scalability and Generalization: Methods that do explore more complex systems and even physical interaction often fail to scale to dense object environments. Reinforcement learning approaches have been used to handle complex dynamics, either by integrating safety as a soft constraint [16], [17] or by using Bellman-inspired updates [8], [9]. However, these methods typically entangle safety with task-specific policies, limiting their use as general, task-agnostic safety filters. Other approaches have leveraged advances in simulation technology such as XPBD [18] and IsaacLab [19] to enable planning for complex systems. For instance, [6] uses XPBD with model based planning to manipulate deformables under safety constraints for active sensing. Other approaches like [7] leverage learned latent space dynamics to reason about more complex definitions of safety in deformable manipulation tasks. Methods like [20] even incorporate additional sensing modalities, like force feedback to safely manipulate a single articulated object.

While effective for single-object settings, these methods remain computationally intensive or fail to generalize well to multi-object, cluttered environments. A key requirement is a scalable method that can reason about a variable number of movable objects without being overly conservative.

Control Barrier Functions: To address the challenges presented above, we build upon the framework of Control Barrier Functions (CBFs). A CBF is a function that defines a control invariant safe set in the system's state space, guaranteeing the existence of a control action that keeps the agent in this set [2], [21]. CBFs have been widely used to enforce safety online as minimally invasive safety filters, with extensions such as Robust CBFs (RCBFs) [22] generalizing them to uncertain dynamics.

A key challenge for classical CBFs is their synthesis [23], as it often requires precise system knowledge. This limitation has historically restricted their use to simpler tasks such as collision avoidance. While HJR-based synthesis is possible [10], it suffers from poor scalability. To overcome this, Neural Control Barrier Functions (NCBFs) learn safety constraints directly from data [24], enabling a new class of methods that incorporate perception [25], handle uncertainty [26], and scale to dense environments [27]. Other works continue to build upon this, enabling refining candidate NCBFs with HJR [28]. Works such as [27] have improved the scalability of safety in dense environments, leveraging composable CBFs to enable safe navigation in crowds. Despite this progress, the application of CBFs remains focused on collision avoidance, leaving a critical gap in methods that can manage safe physical interaction.

In this work, we tackle the problem of safe interaction in dense, multi-object environments where physical contact is unavoidable. We extend CBFs beyond collision avoidance by introducing Dense Contact Barrier Functions which reason about per-object safety in a composable manner. This approach yields scalable, generalizable safety filters that handle diverse configurations while maintaining the ability to interact and complete tasks.

III. Preliminaries

A. Problem Statement

We consider the problem of a robot safely operating in environments densely populated with objects, such as a household pantry or the shelves of a store or warehouse. The robot may physically interact with the objects, but must ensure that its actions do not cause the objects to enter an unsafe state. Though our formulation generalizes to a variety of state-based safety constraints, in this paper we focus on the failure condition of objects tipping beyond a prescribed tilt angle. This problem statement captures a variety of relevant-tasks such as a manipulator reaching for a glass in a pantry, a ground robot pushing aside furniture, or warehouse robots reorganizing shelves. In each case, safety is defined in terms of object stability , while the robot must remain effective at task execution.

We begin by formalizing this problem. We define the state of the robot as $\mathbf{r} \in \mathbb{R}^n$ with dynamics

$$\mathbf{r}^{t+1} = f(\mathbf{r}^t, u^t), \tag{1}$$

where $u^t \in \mathbb{R}^a$ is the control input. The environment contains N objects each with a state $\mathbf{o}_i \in \mathbb{R}^m \ \forall i \in [1, N]$. The safety of each object is defined by a function $q_i(\mathbf{o}_i)$, where a violation occurs if this safety cost exceeds a known threshold, i.e. $q_i(\mathbf{o}_i) \geq q_T$.

For a scene with a single object, $\mathbf{o}$, the dynamics are given by: $\mathbf{o}^{(t+1)} = p(\mathbf{r}^t, u^{(t)}, \mathbf{o}^{(t)})$. However for multiple objects the dynamics, given by

$$\begin{aligned}\{\mathbf{o}_1^{(t+1)}, &\mathbf{o}_2^{(t+1)}, \dots, \mathbf{o}_N^{(t+1)}\} \\ &= F(\mathbf{r}^{(t)}, u^{(t)}, \{\mathbf{o}_1^{(t)}, \mathbf{o}_2^{(t)}, \dots, \mathbf{o}_N^{(t)}\}),\end{aligned} \tag{2}$$

are influenced by both the robot and inter-object interactions. A key challenge with reasoning over the interactions in F is that the number of inter-object interactions scales exponentially with the number of objects. Reasoning about the safety of each object through explicit modeling quickly becomes infeasible. Additionally, physics models for contact for dense environments are often too computationally inefficient for use. In this paper we assume access to the forward dynamics model of the robot (1), but do not have runtime access to the physics model (2), that forward simulates the objects.

B. Control Barrier Functions

Consider a discrete-time system with dynamics $\mathbf{x}^{t+1} = f(\mathbf{x}^t, u^t)$ and state space $\mathbf{x} \in \mathbb{R}^n$. We say that a set $\mathcal{C} \subseteq \mathbb{R}^n$ is *control-invariant* if for any $\mathbf{x} \in \mathbb{R}^n$, there is an action u which keeps the subsequent state in the set, i.e. $f(\mathbf{x}, u) \in \mathcal{C}$.

Moreover, we say a function $h : \mathbb{R}^n \to \mathbb{R}$ is a (discrete-time) *control barrier function* (CBF) if for each state $\mathbf{x} \in \mathbb{R}^n$, the condition

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

holds, where $\gamma : \mathbb{R} \to \mathbb{R}$ is any extended class $\mathcal{K}$ function.

We note that the super-level set $\mathcal{C}_h := \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq 0\}$ for a CBF h is control-invariant.

IV. Dense Contact Barrier Functions

We aim to implicitly capture inter-object contact dynamics using a data-driven neural CBF that certifies safety on a per-object basis. The proposed model is conditioned on the robot state expressed in an object-centric frame and an interaction history summarizing recent contact-induced motion. At inference, these per-object certificates are composed into a global safety assessment to generalize to previously unseen scene densities, yielding a task-agnostic safety filter that permits necessary physical contact during navigation without violating safety constraints. This section develops the theoretical foundations of the proposed implicit interaction modeling and details the CBF formulation.

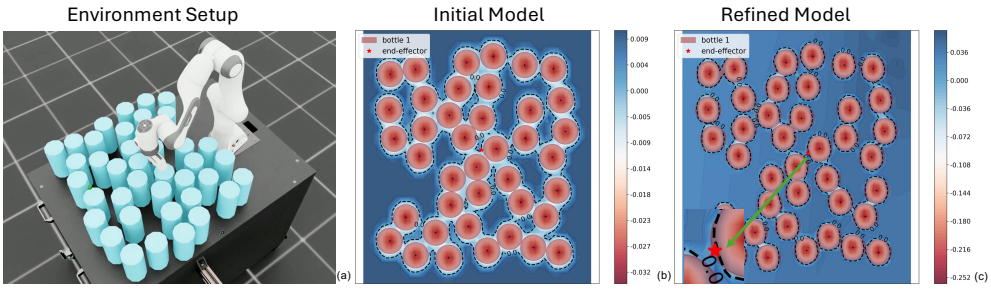

Fig. 2: Comparison of safe boundaries for the initial model and refined model ($\sigma = 0.01$): (a) Environment snapshot with contact. The end-effector makes slight contact, while keeping the object within the safe region. (b) Global CBF value plot for the initial model under contact. The conservative boundary incorrectly treats the interaction as unsafe, leading to the robot stalling. (c) Global CBF value plot for the refined model under contact. The boundary accommodates safe interactions, reflecting less conservative, realistic reasoning about safety for dense environments. Notably, the boundary is not strictly outside the object geometry. The zoomed-in region shows that the boundary locally extends into the object's geometry, indicating that the model enables physically realistic, safe interactions without violating the tilt threshold.

A. Implicit interaction

We first consider a single-object setting with known robot dynamics (1) and unknown contact-driven object dynamics $\mathbf{o}^{t+1} = p(\mathbf{o}^t, \mathbf{r}^t, u^t)$. As accurate analytical contact models are often unavailable or too expensive to simulate online, we introduce a data-driven nominal model $\hat{p}$ to approximate the true dynamics:

$$p(\mathbf{o}^t, \mathbf{r}^t, u^t) = \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u^t) + e(\mathbf{o}^t, \mathbf{r}^t, u^t), \tag{3}$$

where e represents model error. Let $h(\mathbf{r}^t, \mathbf{o}^t)$ be a discrete-time barrier function over the joint (robot, object) state space, and $\sigma > 0$ an upper bound on the barrier-value discrepancy induced by model mismatch, i.e., for all $(\mathbf{r}^t, \mathbf{o}^t, u^t)$,

$$h\big(f(\mathbf{r}^t, u^t), \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u^t)\big) - h\big(f(\mathbf{r}^t, u^t), p(\mathbf{o}^t, \mathbf{r}^t, u^t)\big) \le \sigma.$$

We say that h is a robust CBF [22] with respect to the nominal dynamics $\hat{p}$ if for every $(\mathbf{r}^t, \mathbf{o}^t)$, there exists a control input $u^t \in \mathbb{R}^a$ satisfying

$$h\big(f(\mathbf{r}^t, u^t), \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u^t)\big) - h(\mathbf{r}^t, \mathbf{o}^t) - \sigma \ge -\gamma(h(\mathbf{r}^t, \mathbf{o}^t)).$$

This condition bounds the one-step decrease of h under the nominal model while reserving σ to accommodate bounded model error. Importantly, any u^t satisfying the robust constraint also satisfies the standard CBF condition for the true dynamics p, since the mismatch bound implies

$$\begin{aligned} h\big(f(\mathbf{r}^t, u), p(\mathbf{o}^t, \mathbf{r}^t, u)\big) + \sigma \;&\ge\; h\big(f(\mathbf{r}^t, u), \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u)\big) \\ &\ge\; h(\mathbf{r}^t, \mathbf{o}^t) + \sigma - \gamma\big(h(\mathbf{r}^t, \mathbf{o}^t)\big), \end{aligned}$$

which yields

$$h\big(f(\mathbf{r}^t, u^t), p(\mathbf{o}^t, \mathbf{r}^t, u^t)\big) - h(\mathbf{r}^t, \mathbf{o}^t) \ge -\gamma(h(\mathbf{r}^t, \mathbf{o}^t)).$$

Consequently, if h is a robust CBF with respect to $\hat{p}$ and the mismatch is bounded by σ, then h is also a CBF with respect to the true dynamics p. Define the admissible set

$$K(\mathbf{r}^t, \mathbf{o}^t) = \left\{ u \in \mathbb{R}^a \;\middle|\; \begin{array}{l} h\big(f(\mathbf{r}^t, u), \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u)\big) \\ \quad - h(\mathbf{r}^t, \mathbf{o}^t) - \sigma \ge -\gamma\big(h(\mathbf{r}^t, \mathbf{o}^t)\big) \end{array} \right\}.$$

Any control $u^t \in K(\mathbf{r}^t, \mathbf{o}^t)$ guarantees safety under p.

A sufficient margin σ can be derived via Lipchitz continuity: if h is Lipchitz in $\mathbf{o}^t$ with constant L_o, then the mean value theorem implies the existence of $\mathbf{o}^*$ satisfying

$$\begin{aligned} &\Big|h\big(f(\mathbf{r}^t, u), p(\mathbf{o}^t, \mathbf{r}^t, u)\big) - h\big(f(\mathbf{r}^t, u), \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u)\big)\Big| \\ &= \Big|\frac{\partial}{\partial \mathbf{o}} h\big(f(\mathbf{r}^t, u^t), \mathbf{o}^*\big) \cdot \big(p(\mathbf{o}^t, \mathbf{r}^t, u) - \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u)\big)\Big| \\ &\le L_o \Big|p(\mathbf{o}^t, \mathbf{r}^t, u) - \hat{p}(\mathbf{o}^t, \mathbf{r}^t, u)\Big| \;= L_o \Big|e(\mathbf{o}^t, \mathbf{r}^t, u)\Big|. \end{aligned}$$

Let $\varepsilon \ge 0$ denote a global bound on model error, i.e., $\|e\| \le \varepsilon$ for all admissible inputs. Then choosing $\sigma^* := L_0 \varepsilon$ is sufficient. As $\hat{p}$ improves (i.e., ε decreases), the required robustness margin shrinks accordingly.

Rather than explicitly learning the object dynamics which are highly contact-dependent and often do not generalize to unseen environments, we learn a safety landscape that captures control effect implicitly. When evaluating a candidate control u at $(\mathbf{r}^t, \mathbf{o}^t)$, the quantity we ideally need is

$$h\big(f(\mathbf{r}^t, u), p(\mathbf{o}^t, \mathbf{r}^t, u)\big) = h(\mathbf{r}^{t+1}, \mathbf{o}^{t+1}),$$

which depends on the unknown dynamics p. Instead, we compute $\mathbf{r}^{t+1} = f(\mathbf{r}^t, u)$ using the known robot model, and score the action via a data-driven predictor $B(\mathbf{r}^{t+1}, \mathbf{o}^t)$ serving as a proxy for the "next-step" barrier score: $B(\mathbf{r}^{t+1}, \mathbf{o}^t) \approx h(\mathbf{r}^{t+1}, \mathbf{o}^{t+1})$. If the prediction mismatch is bounded by the robustness margin:

$$B(\mathbf{r}^{t+1}, \mathbf{o}^t) - h(\mathbf{r}^{t+1}, \mathbf{o}^{t+1}) \le \sigma,$$

then enforcing the robust CBF constraints with B implies the standard CBF condition under the true dynamics:

$$B(\mathbf{r}^{t+1}, \mathbf{o}^t) - B(\mathbf{r}^t, \mathbf{o}^{t-1}) - \sigma \ge -\gamma(B(\mathbf{r}^t, \mathbf{o}^{t-1})). \tag{4}$$

By bypassing explicit dynamics modeling, B focuses on the invariant structure required for safety, rather than reconstructing complex, non-smooth contact interactions.

B. Barrier Formulation

1) Temporal Object-State History: To improve implicit inference of interaction dynamics and unobserved physical

parameters (e.g., mass, friction), we follow [27] and augment the input with a finite history of object states. Specifically, we define an absolute history window of length T as $\mathbf{O}^t = [\mathbf{o}^t, \mathbf{o}^{t-1}, \ldots, \mathbf{o}^{t-(T-1)}]$. This temporal context allows B to condition its safety estimate on recent motion trends rather than relying on a single instantaneous observation.

2) Object-Centric Representation: Directly learning a joint barrier over the full multi-object state is generally intractable due to the combinatorial growth of interaction modes. Inspired by [29], we adopt an object-centric representation to learn a single per-object model that can be deployed across variable-sized object sets. Let $R(a, b)$ denote a transformation that expresses state a relative to reference b. For each object i, we anchor the coordinate frame at the object's position at time $t - T$ and define

$$\mathbf{r}_i^t = R(\mathbf{r}^t, \mathbf{o}_i^{t-T}),$$
$$\mathbf{O}_i^t = \Big[R(\mathbf{o}_i^t, \mathbf{o}_i^{t-T}), \ldots, R(\mathbf{o}_i^{t-(T-1)}, \mathbf{o}_i^{t-T})\Big].$$

In this representation, deviations from the "expected" object motion (e.g., lateral drift when pushed forward) implicitly encode the presence of object–object interactions.

3) Inference for Multi-Object Scenes: To certify scene-level safety, we compose the per-object barrier values via a minimum operator. For N objects, the safety for a candidate control input u is evaluated as:

$$B^u_{\text{global}}(\mathbf{r}^t_{1:N}, \mathbf{O}^t_{1:N}) = \min_{i=1,\ldots,N} B(f(\mathbf{r}_i^t, u), \mathbf{O}_i^t). \quad (5)$$

A control input is deemed admissible if $B_{\text{global}} \geq 0$, which is equivalent to requiring that all object-level barrier values are nonnegative (safe). This composition scales linearly in N and is trivially parallelizable across objects.

V. Control Barrier Training and Deployment

In this section, we describe the training procedure for the proposed Dense CBF and show how it is deployed online as a minimally invasive, sampling-based safety filter. The training process involves three stages: 1) interaction-rich **data collection**, 2) **initial training** of a neural barrier model, and 3) iterative boundary **refinement** to reduce conservativeness and improve invariance. Online execution then uses the refined model to filter a nominal controller.

A. Data Collection

We collect interaction-rich trajectories in simulation using a suboptimal behavior policy. This policy induces both safe and unsafe behaviors, yielding frequent interactions near the safety boundary without requiring a highly tuned controller. Training scenes contain only a finite number of objects, typically far fewer than at deployment.

For each trajectory and each object i, we store transition tuples $\{(\mathbf{r}_i^t, \mathbf{O}_i^{t-1}), (\mathbf{r}_i^{t+1}, \mathbf{O}_i^t)\}$, and assign a binary label using the ground-truth safety criterion at the resulting physical state. The labeled transitions form a safe dataset $\mathcal{D}_s$ and an unsafe dataset $\mathcal{D}_u$, with $\mathcal{D} := \mathcal{D}_s \cup \mathcal{D}_u$. For later refinement, we also store the full simulator state at each transition to enable re-simulation under alternative control inputs.

B. Initial Training

We parameterize B with a neural network comprising (i) an LSTM encoder for the temporal object history, (ii) an MLP encoder for the robot state, and (iii) a fusion head that outputs a scalar barrier value. We obtain an initial model B_θ by minimizing the composite objective $L(\theta)$ defined as:

$$L(\theta) = \eta_s L_s(\theta) + \eta_u L_u(\theta) + \eta_d L_d(\theta), \quad (6)$$
$$L_s(\theta) = \mathbb{E}_{\mathcal{D}_s}\Big[- B_\theta(\mathbf{r}_i^t, \mathbf{O}_i^{t-1})\Big]_+, \quad (7)$$
$$L_u(\theta) = \mathbb{E}_{\mathcal{D}_u}\Big[B_\theta(\mathbf{r}_i^t, \mathbf{O}_i^{t-1})\Big]_+, \quad (8)$$
$$L_d(\theta) = \mathbb{E}_{\mathcal{D}_s \cup \mathcal{D}_u}\Big[(1-\gamma)\, B_\theta(\mathbf{r}_i^t, \mathbf{O}_i^{t-1}) - B_\theta(\mathbf{r}_i^{t+1}, \mathbf{O}_i^t) + \sigma\Big]_+, \quad (9)$$

where $[\cdot]_+$ denotes ReLU, $\eta_s, \eta_u, \eta_d > 0$ are weighting coefficients, and $\gamma \in (0, 1)$ defines the adopted linear class-$\mathcal{K}$ function $\gamma(s) = \gamma s$. The terms L_s and L_u enforce correct sign on labeled samples, while L_d penalizes violations of the discrete-time barrier condition.

C. Refinement

The initial model, trained on trajectories generated by a suboptimal behavior policy, is typically conservative. Because labels in $\mathcal{D}$ reflect the behavior policy's outcomes, many feasible contact-rich states that admit safe actions are nonetheless observed (and labeled) unsafe, introducing false negatives and shrinking the learned safe set. Meanwhile, poor initial training can produce false safe classifications that violate invariance. We address both issues with iterative refinement near the learned boundary.

To identify near-boundary samples, we select transitions from $\mathcal{D}$ such that $|B(\mathbf{r}_i^t, \mathbf{O}_i^{t-1})| \leq \delta$ for a tunable threshold $\delta > 0$. For each selected transition, we re-initialize the full multi-object simulator state and execute a short rollout (up to s steps) under a "safest-action" policy that maximizes the next-step global barrier value:

$$u_{\text{refine}}(\mathbf{r}^t_{1:N}, \mathbf{O}^t_{1:N}) = \arg\max_{u \in \mathbb{R}^a} B^u_{\text{global}}(\mathbf{r}^t_{1:N}, \mathbf{O}^t_{1:N}) \quad (10)$$

All transitions collected during these rollouts are annotated using the ground-truth safety criterion and merged into the training set, after which we fine-tune the model via (6). We repeat this loop until performance stabilizes.

D. Deployment

At runtime, the refined barrier is used as a safety filter that minimally perturbs a nominal control input u_{norm} to ensure safety. We use sampling-based action selection as follows:

1) Evaluate the nominal control. If $B^{u_{\text{norm}}}_{\text{global}}(\mathbf{r}^t_{1:N}, \mathbf{O}^t_{1:N}) \geq 0$, execute u_{norm}.
2) Otherwise, sample a candidate set $\mathbb{U}$ around u_{norm} (e.g., Gaussian perturbations in control space).
3) Form the safe subset $\mathbb{U}_s = \{u \in \mathbb{U} \mid B^u_{\text{global}}(\mathbf{r}^t_{1:N}, \mathbf{O}^t_{1:N}) \geq 0$ and execute $u_{\text{filtered}} = \arg\min_{u \in \mathbb{U}_s} \|u - u_{\text{norm}}\|$.

Evaluating B for all objects at every time step can be costly. In practice, it is often sufficient to only evaluate objects that

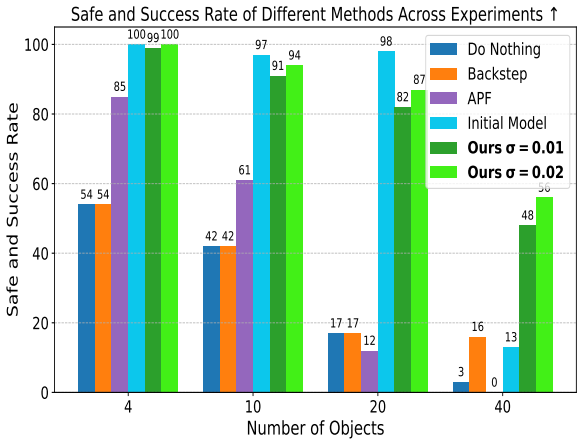

Fig. 3: Safe and success rate of different methods evaluated across environments with increasing object density (4, 10, 20, and 40 objects). This metric illustrates percentage of trajectories where the robot both reached the goal and stayed safe throughout the trajectory. Results are computed over 100 randomized trajectories for each method in each environment. The results highlight the limitations of the baselines, especially at higher object densities and the robustness of our method in maintaining both safety and task success.

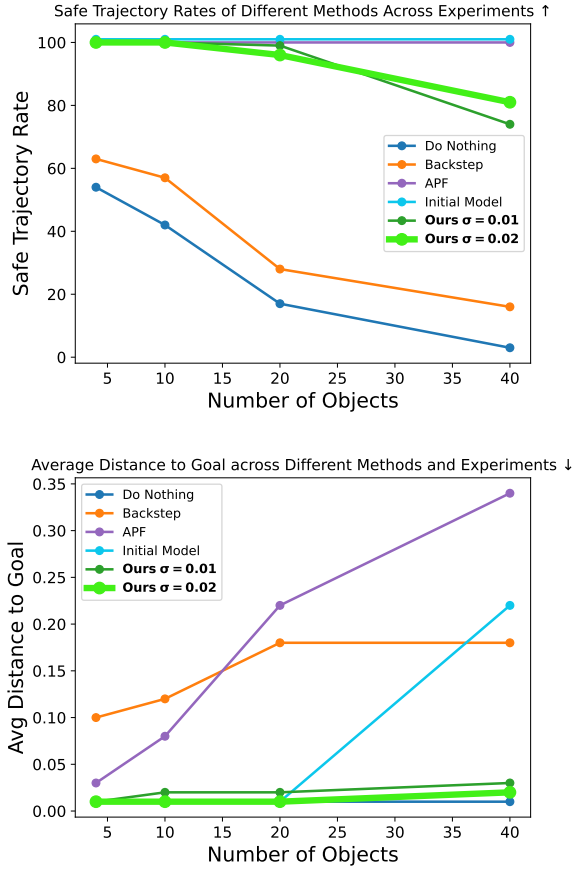

Fig. 4: Performance of different methods in environments of varying object densities (4, 10, 20 and 40 objects). The Safe Rate measures the percentage of safe trajectories (higher is better), and the Average Final Distance to Goal measures final distance to the target (lower is better). These metrics were computed over the same experiment shown in Fig. 3, and provide further insight into the performance of each method. These results highlight the limitations of baselines (*Do Nothing*, *Backstepping*) in staying safe, the conservativeness of *APF*, the conservativeness of the *Initial Model* at higher densities, and the robustness of our proposed method, balancing safety and task success.

are (i) within a fixed spatial neighborhood of the robot or (ii) exhibit recent motion within their history window, as they are likely to be affected by direct or indirect interactions.

VI. Experiments and Results

We evaluate our method in simulation and on hardware to demonstrate safe, scalable contact-rich interaction in cluttered environments where physical contact is unavoidable, and to assess generalization from sparse training scenes to substantially denser clutter at deployment.

Experiment Setup: Experiments are conducted in a custom IsaacLab environment with a Franka Emika Panda arm and cylindrical bottles (Fig. 5). The robot must reach a randomly sampled planar target while ensuring that no bottle exceeds a tilt angle of $15°$. The end-effector is constrained to planar (x, y) motion, and the action space consists of 2D waypoint increments with a maximum step size of 1cm. Each object state is $\mathbf{o}^t = (x^t, y^t, z^t, \theta^t)$, where (x^t, y^t, z^t) denotes the bottle's Cartesian position, and θ^t is the tilting angle. Relative transforms are computed by subtracting (x, y) coordinates in the horizontal plane.

To assess robustness to object heterogeneity, we randomize physical parameters across trials, including mass (1.3–2.0kg) and friction coefficients (static: 0.5–0.7; dynamic: 0.3–0.49). These parameters are not provided as explicit inputs to the learned CBF. Instead, we show that the model can implicitly infer object properties from the temporal history of object-centric states when evaluating candidate controls.

Barrier Training: Models are trained in environments containing only 4 bottles. The initial dataset comprises $1,200$ trajectories collected with a naive back-stepping policy: the robot moves toward a sampled goal, and when any object's tilt approaches the safety threshold, it backtracks to the previous waypoint and resamples a new goal. To increase interaction richness, we balance the dataset by discarding samples in free space.

Next, we report the training parameters.

a) Network: Each timestep contains 4 features: tilt angle and relative object position $(\text{tilt}, \text{rel}_x, \text{rel}_y, \text{rel}_z)$, yielding an input of shape $(B, H, 4)$. An LSTM (hidden size 64) encodes the history sequence, and the last hidden state $(B, 64)$ serves as the object representation. This is concatenated with the planar end-effector position $(x, y) \in \mathbb{R}^2$, forming a 66-dimensional vector. A two-layer MLP (hidden size 64) outputs a scalar CBF value.

b) Training Parameters: We use history $H = 3$ and train for 60 epochs with batch size 256 and learning rate 10^{-4}, using weight decay $\lambda = 0.01$. The CBF parameters

Fig. 5: Method comparison in Real robot experiment. The first row shows the *Do-Nothing* baseline on the real robot, where the end-effector directly pushes through the clutter, resulting in object being knocked over. The second row shows our refined model, which enables controlled contact and safely navigates through the dense environment. The comparison demonstrates how the refined model allows for contact while preserving safety in real-world settings.

are $\gamma = 0.1$, margin $\epsilon = 0.02$, and derivative threshold $\delta = 0.025$, which is tightened to 0.02 during refinement.

Simulation Evaluation: We first evaluate the proposed models in simulation and stress-test generalization by deploying them in scenes that are substantially denser than those used for training, with 4, 10, 20, and 40 objects.

We compare against four baselines: (i) **Do Nothing**, moves directly toward the goal; (ii) **Back-stepping**, the reactive data-collection policy that retreats when any object's tilt approaches $14°$; (iii) **Artificial Potential Fields (APF)**, a repulsive-field method with parameters well-tuned to allow limited contact without excessive conservativeness; and (iv) **Initial:** the proposed barrier model before refinement. We ablate the effect of the robustness margin by evaluating our method with $\sigma = 0.01$ and $\sigma = 0.02$. When applying our safety filter, Do Nothing serves as the nominal controller.

Quantitatively, Fig. 3 reports the safe-and-success rate, i.e., the fraction of trajectories that both reach the goal and remain safe. Fig. 4 further decomposes performance into the safe rate and the average distance to the goal at the end of each episode. Do Nothing and Back-stepping frequently violate safety even at low object densities, and their performance degrades rapidly as clutter increases, indicating that they do not account for robot–object and object–object interaction effects. APF and the unrefined initial model maintain high safe-trajectory rates even in dense scenes, but primarily by being overly conservative and stalling far from the goal; the proposed initial model is less conservative relative to APF. The refined models achieve a substantially better safety–performance trade-off as clutter increases, preserving safety without excessive conservativeness. Comparing $\sigma = 0.01$ and $\sigma = 0.02$ highlights the robustness–conservativeness trade-off: if σ underestimates the required margin, safety violations may occur, whereas large σ improves robustness at the cost of increased conservativeness. Notably, despite training on only 4-object scenes, the refined model generalizes effectively to 40-object scenes, supporting the scalability of the object-centric compositional design.

TABLE I: Quantitative comparison across two scenarios. Arrows indicate whether higher or lower is better.

Metrics	Methods	Scenario 1 - 5 objects	Scenario 2 - 10 objects
Number of Successful Trajectory ↑	Ours-refined model	10	8
	Ours-initial-model	8	2
	APF	1	0
	Do-nothing	10	10
Number of Safe Trajectory ↑	Ours-refined model	7	9
	Ours-initial-model	8	9
	APF	10	10
	Do-nothing	2	0
Number of Safe and Successful Trajectory ↑	Ours-refined model	7	7
	Ours-initial-model	6	1
	APF	1	0
	Do-nothing	2	0
Average Final Distance to Goal (cm) ↓	Ours-refined model	1.93	3.22
	Ours-initial model	2.57	15.48
	APF	19.4	23.3
	Do-nothing	1.64	1.52

Fig. 5(b-c) visualize the global barrier induced by the proposed method. The initial model is overly conservative and produces large "blocked" regions that cause the robot to stall. After refinement, the boundary is tighter and admits safe contact-rich pathways that enable goal reaching.

Hardware Evaluation: We evaluate the proposed model zero-shot on hardware after training entirely in simulation. We employ a Franka robot using two Azure Kinect cameras positioned to mitigate occlusions from the arm and end-effector. Object poses are estimated using ArUco markers mounted on top of the objects. We transform the detected marker poses to be consistent with the pose on the object in simulation. Each object has a radius of 7cm and a height of 23cm with a relatively light weight of approximately 44g (varying slightly between objects). Due to the limited field of view of the Kinect sensors, we evaluate on hardware with 5 and 10 objects, arranged with inter-object spacings of approximately 1–3cm. This yields an object density comparable to the 40-object simulation setting. In real-world trials, we define safety as the absence of knocked-over objects

rather than enforcing a strict tilt-angle threshold. Due to the noisiness of the perception pipeline, we instead adopt a binary safety metric based on whether a bottle falls over. Table I reports the hardware results over 10 trials each in the 5-object and 10-object settings. The observed performance mirrors the trends in simulation. The Do Nothing baseline reaches the goal reliably but is rarely safe, while APF is highly conservative at both object densities. The Initial (unrefined) model performs relatively well with 5 objects compared to APF, but degrades substantially in the 10-object setting. The refined model provides the best overall safety–performance trade-off across both object densities: it attains the highest safe-and-success rate and a low average final distance to the goal, while keeping the average number of knocked-over objects per trial low.

VII. Conclusion

In this work, we introduced Dense Contact Barrier Functions, a scalable framework that enables safe robot interaction in densely cluttered environments. Our approach successfully generalizes from training on a few objects to complex scenes with many objects, without requiring explicit modeling of intractable multi-object dynamics. While our results are promising, several exciting directions for future work remain. We plan to validate our method's ability to handle more diverse objects, including those with varying observable parameters. Furthermore, we aim to extend our method to more complex scenarios involving unconstrained, full-body robot motion, deformable manipulation and, ultimately, demonstrate its effectiveness on physical hardware. This research marks a critical step toward deploying robots that can safely and competently maneuver through the unstructured, contact-rich environments of the real world.

References

[1] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE Press, 2017, p. 2242–2253. [Online]. Available: https://doi.org/10.1109/CDC.2017.8263977

[2] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.

[3] M. C. Lee and M. G. Park, "Artificial potential field based path planning for mobile robots using a virtual obstacle concept," in *Proceedings 2003 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM 2003)*, vol. 2, 2003, pp. 735–740 vol.2.

[4] Y. Tang, X. Chu, J. Huang, and K. W. Samuel Au, "Learning-based mpc with safety filter for constrained deformable linear object manipulation," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2877–2884, 2024.

[5] X. Ding, H. Wang, Y. Ren, Y. Zheng, C. Chen, and J. He, "Safety-critical optimal control for robotic manipulators in a cluttered environment," 11 2022.

[6] N. Shinde, X. Liang, F. Liu, Y. Zhang, F. Richter, S. Herbert, and M. Yip, "Jiggle: An active sensing framework for boundary parameters estimation in deformable surgical environments," 07 2024.

[7] K. Nakamura, L. Peters, and A. Bajcsy, "Generalizing safety beyond collision-avoidance via latent-space reachability analysis," 2025. [Online]. Available: https://arxiv.org/abs/2502.00935

[8] O. So, C. Ge, and C. Fan, "Solving minimum-cost reach avoid using reinforcement learning," 2024. [Online]. Available: https://arxiv.org/abs/2410.22600

[9] W. Sharpless, D. Hirsch, S. Tonkens, N. Shinde, and S. Herbert, "Dual-objective reinforcement learning with novel hamilton-jacobi-bellman formulations," 2025. [Online]. Available: https://arxiv.org/abs/2506.16016

[10] A. Begzadic, N. Shinde, S. Tonkens, D. Hirsch, K. Ugalde, M. Yip, J. Cortes, and S. Herbert, "Back to base: Towards hands-off learning via safe resets with reach-avoid safety filters," in *Proceedings of the 7th Annual Learning for Dynamics & Control Conference*, ser. Proceedings of Machine Learning Research, N. Ozay, L. Balzano, D. Panagou, and A. Abate, Eds., vol. 283. PMLR, 04–06 Jun 2025, pp. 1154–1166. [Online]. Available: https://proceedings.mlr.press/v283/begzadic25a.html

[11] S. Bansal and C. Tomlin, "Deepreach: A deep learning approach to high-dimensional reachability," 2020. [Online]. Available: https://arxiv.org/abs/2011.02082

[12] Z. Feng, L. Qiu, and S. Bansal, "Bridging model predictive control and deep learning for scalable reachability analysis," 2025. [Online]. Available: https://arxiv.org/abs/2505.03830

[13] D. Morton and M. Pavone, "Safe, task-consistent manipulation with operational space control barrier functions," 2025. [Online]. Available: https://arxiv.org/abs/2503.06736

[14] D. A. Dsouza, S. Shenoy, M. Wang, and A. R. Chowdhury, "A comprehensive safety architecture for human–robot collaboration in confined workspaces using improved artificial potential field," *Robotica*, vol. 43, no. 4, pp. 1373–1393, apr 2025.

[15] J. Lee, M. Seo, A. Bylard, R. Sun, and L. Sentis, "Real-time model predictive control for industrial manipulators with singularity-tolerant hierarchical task control," 2022.

[16] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," 2017. [Online]. Available: https://arxiv.org/abs/1705.10528

[17] S. Gu, B. Sel, Y. Ding, L. Wang, Q. Lin, M. Jin, and A. Knoll, "Balance reward and safety optimization for safe reinforcement learning: A perspective of gradient manipulation," 2025. [Online]. Available: https://arxiv.org/abs/2405.01677

[18] M. Macklin, M. Müller, and N. Chentanez, "Xpbd: position-based simulation of compliant constrained dynamics," in *Proceedings of the 9th International Conference on Motion in Games*, 2016, pp. 49–54.

[19] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg, "Orbit: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023.

[20] L. Wei, J. Ma, Y. Hu, and R. Zhang, "Ensuring force safety in vision-guided robotic manipulation via implicit tactile calibration," 2024.

[21] E. Shakhesi, A. Katriniok, and W. P. M. H. M. Heemels, "Counterexample-guided synthesis of robust discrete-time control barrier functions," *IEEE Control Systems Letters*, vol. 9, p. 1574–1579, 2025. [Online]. Available: http://dx.doi.org/10.1109/LCSYS.2025.3578971

[22] A. Alan, T. G. Molnar, A. D. Ames, and G. Orosz, "Parameterized barrier functions to guarantee safety under uncertainty," *IEEE Control Systems Letters*, vol. 7, pp. 2077–2082, 2023.

[23] K. Garg, J. Usevitch, J. Breeden, M. Black, D. Agrawal, H. Parwana, and D. Panagou, "Advances in the theory of control barrier functions: Addressing practical challenges in safe control synthesis for autonomous and robotic systems," 2023. [Online]. Available: https://arxiv.org/abs/2312.16719

[24] S. Liu, C. Liu, and J. Dolan, "Safe control under input limits with neural control barrier functions," 2022. [Online]. Available: https://arxiv.org/abs/2211.11056

[25] S. Raja, A. Kumar, and P. Gupta, "Ogm-cbf: Occupancy-grid map based control barrier functions for safe robot control," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2024, pp. 5678–5685.

[26] M. Tayal, H. Zhang, P. Jagtap, and A. Clark, "Learning a formally verified control barrier function in stochastic environment," 2024.

[27] H. Yu, C. Hirayama, C. Yu, S. Herbert, and S. Gao, "Sequential neural barriers for scalable dynamic obstacle avoidance," 2023. [Online]. Available: https://arxiv.org/abs/2307.03015

[28] S. Tonkens and S. Herbert, "Refining control barrier functions through hamilton-jacobi reachability," 2022. [Online]. Available: https://arxiv.org/abs/2204.12507

[29] N. U. Shinde, J. Johnson, S. Herbert, and M. C. Yip, "Object-centric representations for interactive online learning with non-parametric methods," 2023. [Online]. Available: https://arxiv.org/abs/2307.10063